


Handbook of
THE PHILOSOPHY OF SCIENCE

General Editors: DOV M. GABBAY, PAUL THAGARD, AND JOHN WOODS

PHILOSOPHY
of PHYSICS

PART A



Edited by Jeremy Butterfield
and John Earman



Philosophy of Physics

Part A

Handbook of the Philosophy of Science

General Editors

Dov M. Gabbay
Paul Thagard
John Woods

Cover: Photograph of Albert Einstein taken in 1948 at Princeton
by Yousuf Karsh



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON • NEW YORK • OXFORD
PARIS • SAN DIEGO • SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

North-Holland is an imprint of Elsevier



Philosophy of Physics

Part A

Edited by

Jeremy Butterfield

All Souls College, University of Oxford, Oxford, UK

and

John Earman

Department of History and Philosophy of Science,
University of Pittsburgh, Pittsburgh, PA, USA



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON • NEW YORK • OXFORD
PARIS • SAN DIEGO • SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

North-Holland is an imprint of Elsevier



North-Holland is an imprint of Elsevier
Radarweg 29, PO Box 211, 1000 AE Amsterdam, The Netherlands
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, UK

First edition 2007

Copyright © 2007 Elsevier B.V. All rights reserved

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means electronic, mechanical, photocopying, recording or otherwise without the prior written permission of the publisher

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone (+44) (0) 1865 843830; fax (+44) (0) 1865 853333; email: permissions@elsevier.com. Alternatively you can submit your request online by visiting the Elsevier web site at <http://elsevier.com/locate/permissions>, and selecting *Obtaining permission to use Elsevier material*

Notice

No responsibility is assumed by the publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made

Library of Congress Cataloging-in-Publication Data

A catalog record for this book is available from the Library of Congress

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN-13 (Set): 978-0-444-51560-5
ISBN-10 (Set): 0-444-51560-7
ISBN-13 (Part A): 978-0-444-53001-1
ISBN-10 (Part A): 0-444-53001-0
ISBN-13 (Part B): 978-0-444-53002-8
ISBN-10 (Part B): 0-444-53002-9

For information on all North-Holland publications visit our website at books.elsevier.com
--

Printed and bound in The Netherlands

07 08 09 10 11 10 9 8 7 6 5 4 3 2 1

GENERAL PREFACE

Dov Gabbay, Paul Thagard, and John Woods

Whenever science operates at the cutting edge of what is known, it invariably runs into philosophical issues about the nature of knowledge and reality. Scientific controversies raise such questions as the relation of theory and experiment, the nature of explanation, and the extent to which science can approximate to the truth. Within particular sciences, special concerns arise about what exists and how it can be known, for example in physics about the nature of space and time, and in psychology about the nature of consciousness. Hence the philosophy of science is an essential part of the scientific investigation of the world.

In recent decades, philosophy of science has become an increasingly central part of philosophy in general. Although there are still philosophers who think that theories of knowledge and reality can be developed by pure reflection, much current philosophical work finds it necessary and valuable to take into account relevant scientific findings. For example, the philosophy of mind is now closely tied to empirical psychology, and political theory often intersects with economics. Thus philosophy of science provides a valuable bridge between philosophical and scientific inquiry.

More and more, the philosophy of science concerns itself not just with general issues about the nature and validity of science, but especially with particular issues that arise in specific sciences. Accordingly, we have organized this Handbook into many volumes reflecting the full range of current research in the philosophy of science. We invited volume editors who are fully involved in the specific sciences, and are delighted that they have solicited contributions by scientifically-informed philosophers and (in a few cases) philosophically-informed scientists. The result is the most comprehensive review ever provided of the philosophy of science.

Here are the volumes in the Handbook:

Philosophy of Science: Focal Issues, edited by Theo Kuipers.

Philosophy of Physics, edited by Jeremy Butterfield and John Earman.

Philosophy of Biology, edited by Mohan Matthen and Christopher Stephens.

Philosophy of Mathematics, edited by Andrew Irvine.

Philosophy of Logic, edited by Dale Jacquette.

Philosophy of Chemistry and Pharmacology, edited by Andrea Woody and Robin Hendry.

Philosophy of Statistics, edited by Prasanta S. Bandyopadhyay and Malcolm Forster.

Philosophy of Information, edited by Pieter Adriaans and Johan van Benthem.

Philosophy of Technological Sciences, edited by Anthonie Meijers.

Philosophy of Complex Systems, edited by Cliff Hooker and John Collier.

Philosophy of Earth Systems Science, edited by Bryson Brown and Kent Peacock.

Philosophy of Psychology and Cognitive Science, edited by Paul Thagard.

Philosophy of Economics, edited by Uskali Mäki.

Philosophy of Linguistics, edited by Martin Stokhof and Jeroen Groenendijk.

Philosophy of Anthropology and Sociology, edited by Stephen Turner and Mark Risjord.

Philosophy of Medicine, edited by Fred Gifford.

Details about the contents and publishing schedule of the volumes can be found at <http://www.johnwoods.ca/HPS/>.

As general editors, we are extremely grateful to the volume editors for arranging such a distinguished array of contributors and for managing their contributions. Production of these volumes has been a huge enterprise, and our warmest thanks go to Jane Spurr and Carol Woods for putting them together. Thanks also to Andy Deelen and Arjen Sevenster at Elsevier for their support and direction.

ACKNOWLEDGEMENTS

It is a pleasure to thank several people. First and foremost, we are grateful to John Woods, Dawn Collins, Jane Spurr and their colleagues at Lethbridge, Alberta, and London, England for their unflagging support and encouragement. We also thank Brian Hepburn for invaluable help with references and the preparation of the text.

This page intentionally left blank

CONTENTS

General Preface	v
Dov Gabbay, Paul Thagard, and John Woods	
Introduction	xiii
Jeremy Butterfield and John Earman	
List of Contributors	xxiii
Part A	
On Symplectic Reduction in Classical Mechanics	1
Jeremy Butterfield	
The Representation of Time and Change in Mechanics	133
Gordon Belot	
Classical Relativity Theory	229
David B. Malament	
Non-Relativistic Quantum Mechanics	275
Michael Dickson	
Between Classical and Quantum	417
N. P. Landsman	
Quantum Information and Computing	555
Jeffrey Bub	
The Conceptual Basis of Quantum Field Theory	661
Gerard 't Hooft	
Algebraic Quantum Field Theory	731
Hans Halvorson (with an Appendix by Michael Müger)	
Index to Part A and Part B	I-1

Part B

Compendium of the Foundations of Classical Statistical Physics	923
Jos Uffink	
Quantum Statistical Physics	1075
G�rard Emch	
Issues in the Philosophy of Cosmology	1183
George F.R. Ellis	
Quantum Gravity	1287
Carlo Rovelli	
Symmetries and Invariances in Classical Physics	1331
Katherine Brading and Elena Castellani	
Aspects of Determinism in Modern Physics	1369
John Earman	
Index to Part A and Part B	I-1

This volume is dedicated to the memories of

Robert Clifton (1964–2002)

and

James Cushing (1937–2002).

This page intentionally left blank

INTRODUCTION

Jeremy Butterfield and John Earman

1 THE PHILOSOPHY OF PHYSICS TODAY

In the last forty years, philosophy of physics has become a large and vigorous branch of philosophy, and so has amply won its place in a series of Handbooks in the philosophy of science. The reasons for its vigour are not far to seek. As we see matters, there are two main reasons; the first relates to the formative years of analytic philosophy of science, and the second to the last forty years.

First, physics had an enormous influence on the early phase of the analytic movement in philosophy. This influence does not just reflect the fact that for the logical positivists and logical empiricists, and for others such as Russell, physics represented a paradigm of empirical knowledge. There are also much more specific influences. Each of the three main pillars of modern physics — thermal physics, quantum theory and relativity — contributed specific ideas and arguments to philosophical debate. Among the more obvious influences are the following.

Thermal physics and the scientific controversy about the existence of atoms bore upon the philosophical debate between realism and instrumentalism; and the rise of statistical mechanics fuelled the philosophy of probability. As to quantum theory, its most pervasive influence in philosophy has undoubtedly been to make philosophers accept that a fundamental physical theory could be indeterministic. But this influence is questionable since, as every philosopher of science knows (or should know!), indeterminism only enters at the most controversial point of quantum theory: viz., the alleged “collapse of the wave packet”. In any case, the obscurity of the interpretation of quantum theory threw not only philosophers, but also the giants of physics, such as Einstein and Bohr, into vigorous debate: and not only about determinism, but also about other philosophical fundamentals, such as the nature of objectivity. Finally, relativity theory, both special and general, revolutionized the philosophy of space and time, in particular by threatening neo-Kantian doctrines about the nature of geometry.

These influences meant that when the analytic movement became dominant in anglophone philosophy, the interpretation of modern physics was established as a prominent theme in its sub-discipline, philosophy of science. Accordingly, as philosophy has grown, so has the philosophy of physics.

But from the 1960s onwards, philosophy of physics has also grown for a reason external to philosophy. Namely, within physics itself there has been considerable

interest in foundational issues, with results that have many suggestive repercussions for philosophy. Again, there have been various developments within physics, and thereby various influences on philosophy. The result, we believe, is that nowadays foundational issues in the fundamental physical theories provide the most interesting and important problems in the philosophy of physics. We have chosen the topics for this volume in accord with this conviction. In the next Subsection, we will articulate some of these foundational issues, and thereby introduce the Chapters of the volume.

2 CURRENT FOUNDATIONAL ISSUES IN PHYSICS

We will first discuss these issues under five headings. The first three correspond to the three pillars of modern physics mentioned in Section 2.1; i.e. thermal physics, quantum theory and relativity theory. The fourth and fifth concern combinations of these pillars; and lead to speculations about the future of physics. These five headings will provide a way of introducing most of this volume's Chapters, albeit not in the order in which they occur. Then, after these five headings, we will introduce the volume's remaining two Chapters.

2.1 *Thermal physics*

Controversies about the foundations of thermal physics, especially the characterization of the approach to equilibrium, have continued unabated since the days of the field's founding fathers, such as Maxwell and Boltzmann. Some aspects of the original controversies can be seen again in modern discussions. But the controversies have also been transformed by the development of several scientific fields; especially the following three, which have grown enormously since the 1960s:

- (i) classical mechanics, and its offspring such as ergodic theory and chaos theory;
- (ii) quantum thermal physics; and
- (iii) cosmology, which nowadays provides a very detailed and so fruitful context for developing and evaluating Boltzmann's bold idea that the ultimate origin of the "arrow of time" is cosmological.

In this volume, the foundations of thermal physics is represented by the Chapters by Uffink and by Emch, who cover classical and quantum aspects, respectively. Among the topics Uffink discusses, two receive special attention: the evolution of Boltzmann's views, and the mathematical framework of stochastic dynamics. Emch adopts the formalism of algebraic quantum statistical mechanics, and reviews many results about that formalism's notion of equilibrium, i.e. KMS states. Two other Chapters also provide a little stage-setting for Uffink and Emch, though without pursuing the relation to thermal physics: viz. the Chapters by Butterfield on classical mechanics, and by Ellis on cosmology.

2.2 *Quantum theory*

Since the 1960s, the physics community has witnessed a revival of the debates about the interpretation of quantum theory that raged among the theory's founding fathers. In the general physics community, the single most influential author has no doubt been John Bell, not only through his non-locality theorem and the many experiments it engendered, but also through his critique of the "Copenhagen orthodoxy" and his sympathy towards the pilot-wave and dynamical collapse heterodoxies. But in more specialist communities, there have been other crucial factors that have animated the debate. Mathematical physicists have developed a deep understanding of the various relations between quantum and classical theories. Since the 1970s, there has been progress in understanding decoherence, so that nowadays, almost all would accept that it plays a crucial role in the emergence of the classical world from quantum theory. And since the 1990s, the burgeoning fields of quantum information and computation have grown out of the interpretative debates, especially the analysis of quantum non-locality.

In this volume, these topics are taken up by Dickson, Landsman and Bub. Dickson surveys the formalism of non-relativistic quantum theory, and some of the main interpretative issues, including empirical content, quantum uncertainty, the measurement problem, and non-locality. For the most part, Landsman reviews from the perspective of mathematical physics the relations between quantum and classical theories. In particular, he discusses various approaches to quantization and the rigorous treatments of the classical limits $\hbar \rightarrow 0$ and $N \rightarrow \infty$. But Landsman also includes discussions of the Copenhagen interpretation and decoherence. Finally, Bub presents some central ideas and results about quantum information and quantum computation. As a backdrop to this, he also briefly reviews classical information and computation; and he ends by proposing some provocative morals about the interpretation of quantum theory.

2.3 *Relativity theory*

The decades since the 1960s have seen spectacular developments, for both theory and experiment, in general relativity and cosmology. But this Renaissance has also been very fruitful as regards foundational and philosophical issues. Mathematical relativists have continued to deepen our understanding of the foundations of general relativity: foundations which, as mentioned in Section 1, were recognized already in the 1920s as crucial for the philosophy of space and time. And the recent transformation of cosmology from a largely speculative enterprise into a genuine science has both brought various philosophical questions closer to scientific resolution, and made other philosophical questions, e.g. about method and explanation in cosmology, much more pressing.

In this volume, these topics are represented by the Chapters by Malament, Belot and Ellis. Malament first expounds classical relativity. Then he discusses three special topics: the definition of simultaneity in special relativity, the geometrization of Newtonian gravity, and the extent to which causal structure determines

spacetime geometry. Belot’s main aim is to give a clear statement of the “problem of time” as it occurs in classical general relativity; and to do that, he first reviews the way time is represented in simpler classical theories, including mechanics. (Belot’s Chapter thereby complements Butterfield’s: both expound aspects of classical Hamiltonian theories, and stress how some of these aspects reappear in quantum theories.) Ellis first reviews the present state of relativistic cosmological theory and its observational basis; and then investigates nine philosophical themes, including the anthropic principle and the possible existence of multiverses.

So much by way of introducing some foundational issues, and this volume’s corresponding Chapters, arising *within* one of the three pillars: thermal physics, quantum theory and relativity. We turn to issues arising from combining the pillars — or rather, parts of them! We have already adumbrated the combination of the first and second: viz., in quantum thermal physics, reviewed here by Emch. It is the combination of the second and third — quantum theory and relativity — which we must now address. We shall do so under two headings, corresponding to the distinction between special and general relativity. The first corresponds, of course, to quantum field theory, which forms such a deep and well-established framework for particle physics. The second corresponds to the quantum theory of gravity — which unfortunately still remains only a hope and a goal.¹

2.4 Quantum field theory

Although there are relativistic quantum mechanical theories of a fixed number of particles, by far the most important framework combining quantum theory and special relativity is quantum field theory. Broadly speaking, the foundational issues raised by quantum field theory differ from quantum theory’s traditional interpretative issues, about measurement and non-locality (cf. *Quantum theory*, §2.2 above). There are two points here.

- (i) Although quantum field theory of course illustrates the latter issues just as much as elementary quantum theory does, it apparently cannot offer a resolution of them. The measurement problem and the puzzles about non-locality arise so directly from the unitarity and tensor-product features of quantum theories, as to be unaffected by the extra mathematical structures

¹Our image of three pillars prompts the question: what about the combination of thermal physics and relativity? When Einstein’s special theory of relativity won acceptance, the rush was on to revise the various branches of classical physics to make them properly relativistic. In the case of thermodynamics, this program produced disputes about the Lorentz transformation properties of the thermodynamic quantities of heat, temperature and entropy that persisted well into the 1970s; (see [Liu, 1994] for an overview of this debate). As for classical general relativity theory, there does not currently exist a statistical mechanics that incorporates the “gravitational entropy of the universe”, and it seems unlikely that there can be such a theory. But for all anyone knows, the ideas of thermal physics may play a crucial role in the hoped-for quantum theory of gravity. There are hints to that effect from, for example, black hole thermodynamics, the Unruh effect, and Hawking radiation. These topics are discussed briefly in Rovelli’s chapter.

and physical ideas supplied by quantum field theory.² And accordingly, it has seemed to most workers to be wisest to pursue the traditional interpretative issues within non-relativistic quantum theory: if you identify a problem in a simple context, but are confident that it is not an artefact of the context's simplicity, it is surely wisest to attack it there. (And as shown in this volume by Dickson's and Landsman's Chapters, that context is by no means "really simple": non-relativistic quantum theory, and its relation to classical theories, provides an abundance of intricate structure to investigate.)

- (ii) On the other hand, there are several foundational issues that are distinctive of quantum field theory. Perhaps the most obvious ones are: the nature of particles (including the topic of localization), the interpretation of renormalization, the interpretation of gauge structure, and the existence of unitarily equivalent representations of the canonical commutation relations.

In this volume, these topics are taken up by 't Hooft and by Halvorson and Mürger. First, 't Hooft provides an authoritative survey of quantum field theory, from the perspective of particle physics. Among the main topics he expounds are: the quantization of scalar and spinor fields, Feynman path integrals, the ideas of gauge fields and the Higgs mechanism, renormalization, asymptotic freedom and confinement.

Halvorson and Mürger discuss a smaller and distinctively foundational set of issues, using the apparatus of algebraic quantum field theory. (So their use of the algebraic approach complements the uses made by Emch and Landsman.) They discuss the nature of particles and localization, non-locality, the assignment of values to quantities (i.e. the measurement problem) and the definability of quantum fields at individual spacetime points. But they devote most of their effort to the Doplicher-Haag-Roberts theory of superselection. This theory yields deep insights into crucial structures of quantum field theory: in particular, the set of representations, the relation between the field and observable algebras, and gauge groups.

2.5 *Quantum gravity*

Finally, we turn to the combination of quantum theory with general relativity: i.e., the search for a quantum theory of gravity. Here there is of course no established theory, nor even a consensus about the best approach for constructing one. Rather there are various research programmes that often differ in their technical aims, as well as their motivations and conceptual frameworks. In this situation, various

²In some respects relativistic QFT makes the measurement problem worse. In non-relativistic quantum mechanics, the collapse of the state vector is supposed to happen instantaneously; so in the relativistic setting, one would have to develop some appropriate analogue. On the other hand, the modal interpretation of ordinary QM — which arguably provides the best hope for a no-collapse account of quantum measurement — faces formidable obstacles in relativistic quantum field theory; (see [Clifton, 2000] and Halvorson and Mürger, this volume, Section 5).

foundational issues about the “ingredient” theories are cast in a new light. For example, might quantum gravity revoke orthodox quantum theory’s unitarity, and thereby *en passant* solve the measurement problem? And does the general covariance (diffeomorphism invariance) of general relativity represent an important clue about the ultimate quantum nature of space and time?

In this volume, these and related questions are taken up by Rovelli. He also presents details about other topics: for example, the subject’s history, the two main current programmes (string theory and loop quantum gravity), and quantum cosmology. Ellis’ Chapter also discusses quantum cosmology. In this way, and indeed by addressing other fundamental questions about the idea of an “ultimate” physical theory, Ellis’s Chapter provides a natural complement to Rovelli’s.

So much by way of introducing Chapters that correspond to our initial three pillars of modern physics, or to combinations of them. We turn to introducing the volume’s remaining two Chapters. Here our intention has been to provide Chapters whose discussions bridge the divisions between physical theories, and even those between our three pillars. In this connection, it seemed to us that of the various possible themes for such a cross-cutting discussion, the two most appropriate ones were determinism and symmetry.³

Accordingly, Earman discusses how determinism fares in a wide class of theories: his examples range from classical mechanics to proposals for quantum gravity. He also addresses the relations between determinism and other issues: in particular, predictability, the nature of spacetime, and symmetry. Symmetry in classical physics is given a wide-ranging survey by Brading and Castellani. Among other topics, they discuss: Curie’s principle, the advent of group theory into physics, canonical transformation theory, general covariance in general relativity, and Noether’s theorems. Various aspects of symmetry and invariance in quantum physics are discussed in the Chapters by Dickson, Emch, Halvorson, and Landsman. But a synoptic overview of this complex topic remains to be written — which we hope will be taken as a challenge by some of our readers.

Let us sum up this introduction to the Chapters that follow, with two comments that are intended to give the prospective reader — perhaps daunted by the many pages ahead! — some courage.

First, it is obvious that by our lights, there is no sharp line between philosophy of physics and physics itself. So it is no surprise that some of the best work in philosophy of physics is being done by physicists (as witnessed by several contributions to this volume). No surprise: but certainly, to be welcomed. Conversely, to the traditionally trained philosopher, work by philosophers of physics is liable to look more like physics than philosophy. But for us, this blurring of disciplinary boundaries is no cause for concern. On the contrary, it represents an opportunity for philosophy to enrich itself. And in the other direction, philosophers can hope

³Other good candidates include the “direction of time”, or irreversibility, and the constitution of matter. But adding chapters on these or other cross-cutting themes would have made the volume altogether too long.

that the foundations, and even philosophy, of physics can be a source of heuristic ideas for physics. Or at least, physicists' interest in foundational questions now offers philosophers of physics the opportunity of fruitful discussion with physicists.

But agreed: this enrichment of philosophy does not come for free. And the need to master technical material which is often difficult can be a barrier to entering the philosophy of physics. In designing this volume, our response to this problem has of course been, not to try to lower the barrier, at the cost of scholarship and of fostering illusory hopes: rather our strategy has been to commission Chapters that cover their chosen topics as expertly and completely as possible. So to the reader, our message is simple: take heart! Once you are over the barrier, new vistas open for the philosophy of science.

3 OUTLOOK: HALFWAY THROUGH THE WOODS

Finally, we would like to set the stage for this volume, by making two connected comments about the present state of fundamental physics. Though it may seem naive or hubristic for philosophers to make such comments, we believe it is worth the risk. For we think that at the present juncture fundamental physics is unusually open to contributions from philosophical reflection; and it will be clear from our comments that together they represent an invitation to the reader to make such contributions! The first comment concerns the amazing successes of present-day physics; the second, the fact that so much remains to be understood.

3.1 Successes

First, we want to celebrate the extraordinary achievements of modern physics; specifically of quantum theory and relativity theory. We propose to do this by emphasising how contingent, indeed surprising, it is that the basic postulates of relativity and quantum theory have proved to be so successful in domains of application far beyond their originally intended ones.

Examples are legion. We pick out two examples, almost at random. Why should the new chronogeometry introduced by Einstein's special relativity in 1905 for electromagnetism, be extendible to mechanics, thermodynamics and other fields of physics? And why should the quantum theory, devised for systems of atomic dimensions (10^{-8} cm) be good both for scales much smaller (cf. the nuclear radius of ca. 10^{-12} cm) and vastly larger (cf. superconductivity and superfluidity, involving scales up to 10^{-1} cm)? Indeed, much of the history of twentieth century physics is the story of the consolidation of the relativity and quantum revolutions: the story of their basic postulates being successfully applied ever more widely.

The point applies equally well when we look beyond terrestrial physics. We have in mind, first, general relativity. It makes a wonderful story: the theory was created principally by one person, motivated by conceptual, in part genuinely philosophical, considerations — yet it has proved experimentally accurate in all

kinds of astronomical situations. They range from weak gravitational fields such as occur in the solar system — here it famously explains the minuscule portion of the precession of the perihelion of Mercury (43" of arc per century) that was unaccounted for by Newtonian theory; to fields 10,000 times stronger in a distant binary pulsar — which in the last twenty years has given us compelling (albeit indirect) evidence for a phenomenon (gravitational radiation) that was predicted by general relativity and long searched for; and to exotic objects such as black holes. But general relativity is not the only case. Quantum theory has also been extraordinarily successful in application to astronomy: the obvious example is the use of nuclear physics to develop a very accurate and detailed theories of nucleosynthesis in the very early universe, and of stellar structure and evolution.

Indeed, there is a more general point here, going beyond the successes of relativity and quantum theory. Namely, we tend to get used to the various unities in nature that science reveals — and thereby to forget how contingent and surprising they are. Of course, this is not just a tendency of our own era. For example, nineteenth century physics confirmed Newton's law of gravitation to apply outside the solar system, and discovered terrestrial elements to exist in the stars (by spectroscopy): discoveries that were briefly surprising, but soon taken for granted, incorporated into the educated person's 'common sense'. Similarly nowadays: the many and varied successes of physics in the last few decades, in modelling very accurately phenomena that are (i) vastly distant in space and time, and-or (ii) very different from our usual laboratory scales (in their characteristic values of such quantities as energy, temperature, or pressure etc.), reveal an amazing unity in nature. General theoretical examples of such unity, examples that span some 200 years, are: the ubiquitous fruitfulness of the field concept; and more specifically, of least action principles. For a modern, specific (and literally spectacular) example, consider the precision and detail of our models of supernovae; as confirmed by the wonderful capacity of modern telescope technology to see and analyse individual supernovae, even in other galaxies.

3.2 Clouds on the horizon

And yet: complacency, let alone triumphalism, is not in order! Current physics is full of unfinished business — that is always true in human enquiry. But more to the point, there are clouds on the horizon that may prove as great a threat to the continued success of twentieth century physics, as were the anomalies confronting classical physics at the end of the nineteenth century.

Of course, people differ about what problems they find worrisome; and among the worrisome ones, about which problems are now ripe for being solved, or at least worth addressing. As philosophers, we are generalists: so we naturally find all the various foundational issues mentioned above worrisome. But being generalists, we will of course duck out of trying to say which are the closest to solution, or which are most likely to repay being addressed! In any case, such judgments are hard to adjudicate, since intellectual temperament, and the happenstance of what one

knows about or is interested in, play a large part in forming them.

But we would like to end by returning to one of Section 2's "clouds": a cloud which clearly invites philosophical reflection, and perhaps contributions. Namely, the problem of quantum gravity; in other words, the fact that general relativity and quantum theory are yet to be reconciled. As mentioned in Section 2.5, Rovelli (this volume) discusses how the contrasting conceptual structures of the "ingredient" theories and the ongoing controversies about interpreting them, make for conflicting basic approaches to quantum gravity.

But we want here to emphasise another reason why we still lack a successful theory, despite great effort and ingenuity. In short, it is that the successes of relativity and quantum theory, celebrated in Comment 3.1 above, conspire to deprive us of the relevant experimental data.

Thus there are general reasons to expect data characteristic of quantum gravity to arise only in a regime of energies so high (correspondingly, distances and times so short) as to be completely inaccessible to us. To put the point in terms of length: the value of the Planck length which we expect to be characteristic of quantum gravity is around 10^{-33} cm. This is truly minuscule: the diameters of an atom, nucleus, proton and quark are, respectively, about 10^{-8} , 10^{-12} , 10^{-13} , and 10^{-16} cm. So the Planck length is as many orders of magnitude from the (upper limit for) the diameter of a quark, as that diameter is from our familiar scale of a centimetre!

We can now see how quantum gravity research is in a sense the victim of the successes of relativity and quantum theory. For those successes suggest that we will not see any "new physics" intimating quantum gravity even at the highest energies accessible to us. The obvious example is quasars: these are typically a few light-days in diameter, and yet have a luminosity 1000 times that of our galaxy (itself 100,000 light-years across, containing a hundred billion stars). They are the most energetic, distant (and hence past!) celestial objects that we observe: they are now believed to be fuelled by massive black holes in their cores. Yet suggestions, current thirty years ago, that their stupendous energies and other properties that we *can* observe, could only be explained by fundamentally new physics, have nowadays given way to acceptance that "conventional physics" describing events *outside* the black hole's event-horizon can do so. (Agreed, we expect the physics deep inside the black hole, in the vicinity of its singularity, to exhibit quantum gravity effects: but if ever a region deserved the name "inaccessible", this is surely one!) So the situation is ironic, and even frustrating: quantum gravity research is a victim of its ingredient theories' success.

In any case, the search for quantum gravity is wide open. In closing, we would like to endorse an analogy of Rovelli's [1997]. He suggests that our present search is like that of the mechanical philosophers such as Galileo and Kepler of the early seventeenth century. Just as they struggled with the clues given by Copernicus and Brahe, *en route* to the synthesis given by Newton, so also we are "halfway through the woods". Of course we should be wary of too grossly simplifying and periodizing the scientific revolution, and *a fortiori* of facile analogies between different

historical situations. Nevertheless, it is striking what a “mixed bag” the doctrines of figures such as Galileo and Kepler turn out to have been, from the perspective of the later synthesis. For all their genius, they appear to us (endowed with the anachronistic benefits of hindsight), to have been “transitional figures”. One cannot help speculating that to some future reader of twentieth century physics, enlightened by some future synthesis of general relativity and quantum theory, our current and recent efforts in quantum gravity will seem strange: worthy and sensible from the authors’ perspective (one hopes), but a hodge-podge of insight and error from the reader’s!

BIBLIOGRAPHY

- [Clifton, 2000] R. Clifton. The modal interpretation of algebraic quantum field theory. *Physics Letters A*, 271, 167-177, 2000.
- [Liu, 1994] C. Liu. Is there a relativistic thermodynamics? A case study in the meaning of special relativity. *Studies in the History and Philosophy of Science*, 25, 983-1004, 1994.
- [Rovelli, 1997] C. Rovelli. Halfway through the woods. In J. Earman and J. Norton (Eds.), *The Cosmos of Science* (pp. 180-223). Pittsburgh: University of Pittsburgh Press and Konstanz: Universitäts Verlag, 1997.

CONTRIBUTORS

Gordon Belot

Department of Philosophy, University of Pittsburgh, 1001 Cathedral of Learning,
Pittsburgh, PA 15260, USA.

gbelot@pitt.edu

Katherine Brading

Department of Philosophy, University of Notre Dame Notre Dame, 100 Malloy
Hall, Indiana 46556, USA.

kbrading@nd.edu

Jeffrey Bub

Department of Philosophy, University of Maryland, College Park, MD 20742, USA.

jbub@umd.edu

Jeremy Butterfield

Trinity College, Cambridge, CB2 1TQ, UK.

jb56@cus.cam.ac.uk

Elena Castellani

Department of Philosophy, University of Florence, via Bolognese 52, 50139, Firenze,
Italy.

elena.castellani@unifi.it

Michael Dickson

Department of Philosophy, University of South Carolina, Columbia, SC 29208,
USA.

dickson@sc.edu

John Earman

Department of History and Philosophy of Science, 1017 Cathedral of Learning,
University of Pittsburgh, Pittsburgh, PA 15260, USA.

jearman+@pitt.edu

George F. R. Ellis

Mathematics Department, University of Cape Town, Rondebosch, Cape Town
8001, South Africa.

ellis@maths.uct.ac.za

G rard Emch

Department of Mathematics, University of Florida, 358 Little Hall, PO Box 118105, Gainesville, FL 32611–8105, USA.

gge@math.ufl.edu

Hans Halvorson

Department of Philosophy, Princeton University, Princeton, NJ 08544, USA.

hhalvors@princeton.edu

Gerard 't Hooft

Institute for Theoretical Physics, Utrecht University, Leuvenlaan 4, 3584 CC Utrecht, The Netherlands, and Spinoza Institute, Postbox 80.195, 3508 TD Utrecht, The Netherlands.

g.thooft@phys.uu.nl

N.P. Landsman

Institute for Mathematics, Astrophysics, and Particle Physics, Radboud Universiteit Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands

landsman@math.ru.nl

David B. Malament

Department of Logic and Philosophy of Science, University of California at Irvine, 3151 Social Science Plaza, Irvine, CA 92697-5100, USA.

dmalamen@uci.edu

Michael M ger

Institute for Mathematics, Astrophysics, and Particle Physics, Radboud Universiteit Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands.

mueger@math.ru.nl

Carlo Rovelli

Centre de Physique Th orique de Luminy, Universit  de la M diterran e, case 907, F-13288 Marseille, France.

rovelli@cpt.univ-mrs.fr

Jos Uffink

Institute for History and Foundations of Science, Utrecht University, PO Box 80 000, 3508 TA Utrecht, The Netherlands.

uffink@phys.uu.nl

ON SYMPLECTIC REDUCTION IN CLASSICAL MECHANICS

J. Butterfield

Mottoes

The current vitality of mechanics, including the investigation of fundamental questions, is quite remarkable, given its long history and development. This vitality comes about through rich interactions with pure mathematics (from topology and geometry to group representation theory), and through new and exciting applications to areas like control theory. It is perhaps even more remarkable that absolutely fundamental points, such as a clear and unambiguous linking of Lie's work on the Lie-Poisson bracket on the dual of a Lie algebra ... with the most basic of examples in mechanics, such as the rigid body and the motion of ideal fluids, took nearly a century to complete.

Marsden and Ratiu [1999, pp. 431–432]

In the ordinary theory of the rigid body, six different three-dimensional spaces \mathbb{R}^3 , \mathbb{R}^{3*} , \mathfrak{g} , \mathfrak{g}^* , TG_g , T^*G_g are identified.

Arnold [1989, p. 324]

1 INTRODUCTION

1.1 *Why classical mechanics?*

All hail the rise of modern physics! Between 1890 and 1930, the quantum and relativity revolutions and the consolidation of statistical physics through the discovery of atoms, utterly transformed our understanding of nature; and had an enormous influence on philosophy; (e.g. [Kragh, 1999; Ryckman, 2005]). Accordingly, this Handbook concentrates on those three pillars of modern physics — quantum theories, spacetime theories and thermal physics. So some initial explanation of the inclusion of a Chapter on classical mechanics, indeed the classical mechanics of finite-dimensional systems, is in order.

The first point to make is that the various fields of classical physics, such as mechanics and optics, are wonderfully rich and deep, not only in their technicalities, but also in their implications for the philosophy and foundations of physics. From Newton's time onwards, classical mechanics and optics have engendered an enormous amount of philosophical reflection. As regards mechanics, the central

philosophical topics are usually taken (and have traditionally been taken) to be space, time, determinism and the action-at-a-distance nature of Newtonian gravity. Despite their importance, I will not discuss these topics; but some other Chapters will do so (at least in part, and sometimes in connection with theories other than classical mechanics). I will instead focus on the theory of symplectic reduction, which develops the well-known connection between continuous symmetries and conserved quantities, summed up in Noether's "first theorem". I choose this focus partly by way of preparation for parts of some other Chapters; and partly because, as we will see in a moment, symplectic reduction plays a central role in the current renaissance of classical mechanics, and in its relation to quantum physics.

I said that classical physics engendered a lot of philosophical reflection. It is worth stressing two, mutually related, reasons for this: reasons which today's philosophical emphasis on the quantum and relativity revolutions tends to make us forget.

First: in the two centuries following Newton, these fields of classical physics were transformed out of all recognition, so that the framework for philosophical reflection about them also changed. Think of how in the nineteenth century, classical mechanics and optics gave rise to classical field theories, especially electromagnetism. And within this Chapter's specific field, the classical mechanics of finite-dimensional systems, think of how even its central theoretical principles were successively recast, in fundamental ways, by figures such as Euler, Lagrange, Hamilton and Jacobi.

Second, various difficult problems beset the attempt to rigorously formulate classical mechanics and optics; some of which have considerable philosophical aspects. It is *not* true that once we set aside the familiar war-horse topics — space, time, determinism and action-at-a-distance — the world-picture of classical mechanics is straightforward: just "matter in motion". On the contrary. Even if we consider only finite-dimensional systems, we can ask, for example:

- (i) For point-particles (material points): can they have different masses, and if so how? What happens when they collide? Indeed, for point-particles interacting only by Newtonian gravity, a collision involves infinite kinetic energy.
- (ii) For extended bodies treated as finite-dimensional because rigid: what happens when they collide? Rigidity implies that forces, and displacements, are transmitted "infinitely fast" through the body. Surely that should not be taken literally? But if so, what justifies this idealization; and what are its scope and limits?

As to infinite-dimensional systems (elastic solids, fluids and fields), many parts of their theories remain active research areas, especially as regards rigorous formulations and results. For contemporary work on elastic solids, for example, cf. Marsden and Hughes [1982]. As to fluids, the existence and uniqueness of rigorous solutions of the main governing equations, the Navier-Stokes equations, is still an

open problem. This problem not only has an obvious bearing on determinism; it is regarded as scientifically significant enough that its solution would secure a million-dollar Clay Millennium prize.

These two reasons — the successive reformulations of classical mechanics, and its philosophical problems — are of course related. The monumental figures of classical mechanics recognized and debated the problems, and much of their technical work was aimed at solving them. As a result, there was a rich debate about the foundations of classical physics, in particular mechanics, for the two centuries after Newton’s *Principia* [1687]. A well-known example is Duhem’s instrumentalist philosophy of science, which arose in large measure from his realization how hard it was to secure rigorous foundations at the microscopic level for classical mechanics. A similar example is Hilbert’s being prompted by his contemporaries’ continuing controversies about the foundations of mechanics, to choose as the sixth of his famous list of outstanding mathematical problems, the axiomatization of mechanics and probability; (but for some history of this list, cf. [Grattan-Guinness, 2000]). A third example, spanning both centuries, concerns variational principles: the various principles of least action formulated first by Maupertuis, then by Euler and later figures — first for finite classical mechanical systems, then for infinite ones — prompted much discussion of teleology. Indeed, this discussion ensnared the logical empiricists [Stöltzner, 2003]; it also bears on contemporary philosophy of modality [Butterfield, 2004].

In the first half of the twentieth century, the quantum and relativity revolutions tended to distract physicists, and thereby philosophers, from these and similar problems. The excitement of developing the new theories, and of debating their implications for natural philosophy, made it understandable, even inevitable, that the foundational problems of classical mechanics were ignored.

Besides, this tendency was strengthened by the demands of pedagogy: the necessity of including the new theories in physics undergraduate degrees. By mid-century, the constraints of time on the physics curriculum had led many physics undergraduates’ education in classical mechanics to finish with the elementary parts of analytical mechanics, especially of finite-dimensional systems: for example, with the material in Goldstein’s well-known textbook [1950]. Such a restriction is understandable, not least because: (i) the elementary theory of Lagrange’s and Hamilton’s equations requires knowledge of ordinary differential equations, and (ii) elementary Hamiltonian mechanics forms a springboard to learning elementary canonical quantization (as does Hamilton-Jacobi theory, from another perspective). Besides, as I mentioned: even this restricted body of theory provides plenty of material for philosophical analysis — witness my examples above, and the discussions of the great figures such Euler, Lagrange, Hamilton and Jacobi.

However, the second half of the twentieth century saw a renaissance in research in classical mechanics: hence my first motto. There are four obvious reasons for this: the first two “academic”, and the second two “practical”.

- (i) Thanks partly to developments in mathematics in the decades after Hilbert’s list of problems, the foundational questions were addressed afresh, as much

by mathematicians and mathematically-minded engineers as by physicists. The most relevant developments lay in such fields as topology, differential geometry, measure theory and functional analysis. In this revival, the contributions of the Soviet school, always strong in mechanics and probability, were second to none. And relatedly: —

- (ii) The quest to deepen the formulation of quantum theory, especially quantum field theory, prompted investigation of (a) the structure of classical mechanics and (b) quantization. For both (a) and (b), special interest attaches to the generally much harder case of infinite systems.
- (iii) The coming of spaceflight, which spurred the development of celestial mechanics. And relatedly: —
- (iv) The study of non-linear dynamics (“chaos theory”), which was spurred by the invention of computers.

With these diverse causes and aspects, this renaissance continues to flourish — and accordingly, I shall duck out of trying to further adumbrate it! I shall even duck out of trying to survey the philosophical questions that arise from the various formulations of mechanics from Newton to Jacobi and Poincaré. Suffice it to say here that to the various topics mentioned above, one could add, for example, the following two: the first broadly ontological, the second broadly epistemological.

- (a) The analysis of notions such as mass and force (including how they change over time). For this topic, older books include Jammer [1957; 1961] and McMullin [1978]; recent books include [Boudri, 2002; Jammer, 2000; Lutzen, 2005] and [Slovik, 2002]; Grattan-Guinness [2006] is a fine recent synopsis of the history, with many references.
- (b) The analysis of what it is to have an explicit solution of a mechanical problem (including how the notion of explicit solution was gradually generalized). This topic is multi-faceted. It not only relates to the gradual generalization of the notion of function (a grand theme in the history of mathematics — well surveyed by Lutzen [2003]), and to modern non-linear dynamics (cf. (iv) above). It also relates to the simplification of problems by exploiting a symmetry so as to reduce the number of variables one needs — and *this* is the core idea of symplectic reduction. I turn to introducing it.

1.2 *Prospectus*

The strategy of simplifying a mechanical problem by exploiting a symmetry so as to reduce the number of variables is one of classical mechanics’ grand themes. It is theoretically deep, practically important, and recurrent in the history of the subject. The best-known general theorem about the strategy is undoubtedly Noether’s theorem, which describes a correspondence between continuous symmetries and conserved quantities. There is both a Lagrangian and a Hamiltonian

version of this theorem, though for historical reasons the name ‘Noether’s theorem’ is more strongly attached to the Lagrangian version. However, we shall only need the Hamiltonian version of the theorem: it will be the “springboard” for our exposition of symplectic reduction.¹

So I shall begin by briefly reviewing the Hamiltonian version in Section 2.1. For the moment, suffice it to make four comments (in ascending order of importance for what follows):

- (i) Both versions are underpinned by the theorems in elementary Lagrangian and Hamiltonian mechanics about cyclic (ignorable) coordinates and their corresponding conserved momenta.²
- (ii) In fact, the Hamiltonian version of the theorem is stronger. This reflects the fact that the canonical transformations form a “larger” group than the point transformations. A bit more precisely: though the point transformations $q \rightarrow q'$ on the configuration space Q induce canonical transformations on the phase space Γ of the qs and ps , $q \rightarrow q', p \rightarrow p'$, there are yet other canonical transformations which “mix” the qs and ps in ways that transformations induced by point transformations do not.
- (iii) I shall limit our discussion to (a) time-independent Hamiltonians and (b) time-independent transformations. Agreed, analytical mechanics can be developed, in both Lagrangian and Hamiltonian frameworks, while allowing time-dependent dynamics and transformations. For example, in the Lagrangian framework, allowing velocity-dependent potentials and-or time-dependent constraints would prompt one to use what is often called the ‘extended configuration space’ $Q \times \mathbb{R}$. And in the Hamiltonian framework, time-dependence prompts one to use an ‘extended phase space’ $\Gamma \times \mathbb{R}$. Besides, from a philosophical viewpoint, it is important to consider time-dependent transformations: for they include boosts, which are central to the philosophical discussion of spacetime symmetry groups, and especially of relativity

¹For discussion of the Lagrangian version, cf. e.g. Brading and Castellani (this vol., ch. 13) or (restricted to finite-dimensional systems) Butterfield [2004a, Section 4.7]. For an exposition of both versions that is complementary to this paper (and restricted to finite-dimensional systems), cf. Butterfield [2006]. Brading and Castellani also bring out that, even apart from Noether’s theorems in other branches of mathematics, there are other ‘Noether’s theorems’ about symmetries in classical dynamics; so the present theorem is sometimes called Noether’s “first theorem”. Note also (though I shall not develop this point) that symplectic structure can be seen in the classical solution space of Lagrange’s equations, so that symplectic reduction can be developed in the Lagrangian framework; cf. e.g. Marsden and Ratiu [1999, p. 10, Sections 7.2–7.5, and 13.5].

²Here we glimpse the long history of our subject: these theorems were of course clear to these subjects’ founders. Indeed the strategy of exploiting a symmetry to reduce the number of variables occurs already in 1687, in Newton’s solution of the Kepler problem; (or more generally, the problem of two bodies exerting equal and opposite forces along the line between them). The symmetries are translations and rotations, and the corresponding conserved quantities are the linear and angular momenta. In what follows, these symmetries and quantities will provide us with several examples.

principles. But beware: rough-and-ready statements about symmetry, e.g. that the Hamiltonian must be invariant under a symmetry transformation, are liable to stumble on these transformations. To give the simplest example: the Hamiltonian of a free particle is just its kinetic energy, which can be made zero by transforming to the particle's rest frame; i.e. it is not invariant under boosts.

So a full treatment of symmetry in Hamiltonian mechanics, and thereby of symplectic reduction, needs to treat time-dependent transformations — and to beware! But I will set aside all these complications. Here it must suffice to assert, without any details, that the modern theory of symplectic reduction *does* cope with boosts; and more generally, with time-dependent dynamics and transformations.

- (iv) As we shall see in detail, there are three main ways in which the theory of symplectic reduction generalizes Noether's theorem. As one might expect, these three ways are intimately related to one another.
 - (a) Noether's theorem is “one-dimensional” in the sense that for each symmetry (a vector field of a special kind on the phase space), it provides a conserved quantity, i.e. a real-valued function on the phase space, whose value stays constant over time. So in particular, different components of a conserved vector quantity, such as total linear momentum, are treated separately; (in this example, the corresponding vector fields generate translations in three different spatial directions). But in symplectic reduction, the notion of a *momentum map* provides a “unified” description of these different components.
 - (b) Given a symmetry, Noether's theorem enables us to confine our attention to the level surface of the conserved quantity, i.e. the sub-manifold of phase space on which the quantity takes its initial value: for the system's time-evolution is confined to that surface. In that sense, the number of variables we need to consider is reduced. But in symplectic reduction, we go further and form a *quotient space* from the phase space. That is, in the jargon of logic: we define on phase space an equivalence relation (not in general so simple as having a common value for a conserved quantity) and form the set of equivalence classes. In the jargon of group actions: we form the set of orbits. Passage to this quotient space can have various good technical, and even philosophical, motivations. And under good conditions, this set is itself a manifold with lower dimension.
 - (c) Hamiltonian mechanics, and so Noether's theorem, is usually formulated in terms of symplectic manifolds, in particular the cotangent bundle T^*Q of the configuration space Q . (Section 2.1 will give details.) But in symplectic reduction, we often need a (mild) generalization of the idea of a symplectic manifold, called a *Poisson manifold*, in which a

bracket, with some of the properties of the Poisson bracket, is taken as the primitive notion. Besides, this is related to (b) in that we are often led to a Poisson manifold, and dynamics on it, by taking the quotient of a symplectic manifold (i.e. a phase space of the usual kind) by the action of a symmetry group.

As comment (iv) hints, symplectic reduction is a large subject. So there are several motivations for expounding it. As regards physics, many of the ideas and results can be developed for finite-dimensional classical systems (to which I will confine myself), but then generalized to infinite-dimensional systems. And in either the original or the generalized form, they underpin developments in quantum theories. So these ideas and results have formed part of the contemporary renaissance in classical mechanics; cf. (i) and (ii) at the end of Section 1.1.

As regards philosophy, symmetry is both a long-established focus for philosophical discussion, and a currently active one: cf. [Brading and Castellani, 2003]. But philosophical discussion of symplectic reduction seems to have begun only recently, especially in some papers of Belot and Earman. This delay is presumably because the technical material is more sophisticated: indeed, the theory of symplectic reduction was cast in its current general form only in the 1970s. But as Belot and Earman emphasise, the philosophical benefits are worth the price of learning the technicalities. The most obvious issue is that symplectic reduction's device of quotienting a state space casts light on philosophical issues about whether two apparently distinct but utterly indiscernible possibilities should be ruled to be one and the same. In Section 2, I will follow Belot in illustrating this issue with "relationist" mechanics. Indeed, I have selected the topics for my exposition with an eye to giving philosophical readers the background needed for some of Belot's discussions. His papers (which I will cite in Section 2) make many judicious philosophical points, without burdening the reader with an exposition of technicalities: excellent stuff — but to fully appreciate the issues, one of course has to slog through the details.

Finally, in the context of this volume, symplectic reduction provides some background for the Chapters on the representation of time in mechanics [Belot, this vol., ch. 2], and on the relations between classical and quantum physics [Landsman, this vol., ch. 5, especially Sections 4.3–4.5 and 6.5; Dickson, this vol., ch. 4].

The plan of the Chapter is as follows. I first review Noether's theorem in Hamiltonian mechanics as usually formulated, in Section 2.1. Then I introduce the themes mentioned in (b) and (c) above, of quotienting a phase space, and Poisson manifolds (Section 2.2); and illustrate these themes with "relationist" mechanics (Section 2.3).

Thereafter, I expound the basics of symplectic reduction: (confining myself to finite-dimensional Hamiltonian mechanics). Section by Section, the plan will be as follows. Sections 3 and 4 review the modern geometry that will be needed. Section 3 is mostly about Frobenius' theorem, Lie algebras and Lie groups.³ Section

³Its first two Subsections also provide some pre-requisites for Malament (this vol.).

4 expounds Lie group actions. It ends with the central idea of the co-adjoint representation of a Lie group G on the dual \mathfrak{g}^* of its Lie algebra. This review enables us to better understand the motivations for Poisson manifolds (5.1); and then to exhibit examples, and prove some main properties (Section 5.2 onwards). Section 6 applies this material to symmetry and conservation in mechanical systems. In particular, it expresses conserved quantities as momentum maps, and proves Noether’s theorem for Hamiltonian mechanics on Poisson manifolds. Finally, in Section 7, we prove one of the several main theorems about symplectic reduction. It concerns the case where the natural configuration space for a system is itself a Lie group G : this occurs both for the rigid body and ideal fluids. In this case, quotienting the natural phase space (the cotangent bundle on G) gives a Poisson manifold that “is” the dual \mathfrak{g}^* of G ’s Lie algebra.⁴

To sum up: — The overall effect of this exposition is, I hope, to illustrate this Chapter’s mottoes: that classical mechanics is alive and kicking, not least through deepening our understanding of time-honoured systems such as the rigid body — whose analysis in traditional textbooks can be all too confusing!

2 SYMPLECTIC REDUCTION: AN OVERVIEW

We begin by briefly reviewing Hamiltonian mechanics and Noether’s theorem, in Section 2.1.⁵ This prepares us for the idea of symplectic reduction, Section 2.2: which we then illustrate using “relationist” mechanics, Section 2.3.

2.1 *Hamiltonian mechanics and Noether’s theorem: a review*

2.1.1 *Symplectic manifolds; the cotangent bundle as a symplectic manifold*

A *symplectic structure* or *symplectic form* on a manifold M is defined to be a differential 2-form ω on M that is closed (i.e. its exterior derivative $d\omega$ vanishes) and is non-degenerate. That is: for any $x \in M$, and any two tangent vectors at x , $\sigma, \tau \in T_x$:

$$(1) \quad d\omega = 0 \quad \text{and} \quad \forall \tau \neq 0, \exists \sigma : \omega(\tau, \sigma) \neq 0 .$$

Such a pair (M, ω) is called a *symplectic manifold*. There is a rich theory of symplectic manifolds; but we shall only need a small fragment of it. (In particular, the fact that we mostly avoid the theory of canonical transformations means we will not need the theory of Lagrangian sub-manifolds.)

⁴In this endeavour, my sources are four books by masters of the subject: [Abraham and Marsden, 1978; Arnold, 1989; Marsden and Ratiu, 1999; Olver, 2000]. But again, be warned: my selection is severe, as anyone acquainted with these or similar books will recognize.

⁵For more details about differential geometry, cf. Sections 3.1 and 3.2. For more details about the geometric formulation of mechanics, cf. [Arnold, 1989; Marsden and Ratiu, 1999]; or [Singer, 2001] (more elementary than this exposition) or [Abraham and Marsden, 1978] (more advanced); or [Butterfield, 2006] (at the same level). Of many good textbooks of mechanics, I admire especially [Desloge, 1982; Johns, 2005].

First, it follows from the non-degeneracy of ω that M is even-dimensional. The reason lies in a theorem of linear algebra, which one then applies to the tangent space at each point of M . Namely, for any bilinear form $\omega : V \times V \rightarrow \mathbb{R}$: if ω is antisymmetric of rank $r \leq m \equiv \dim(V)$, then r is even. That is: $r = 2n$ for some integer n , and there is a basis $e_1, \dots, e_i, \dots, e_m$ of V for which ω has a simple expansion as wedge-products

$$(2) \quad \omega = \sum_{i=1}^n e^i \wedge e^{i+n} ;$$

equivalently, ω has the $m \times m$ matrix

$$(3) \quad \omega = \begin{pmatrix} \mathbf{0} & \mathbf{1} & \mathbf{0} \\ -\mathbf{1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix} .$$

where $\mathbf{1}$ is the $n \times n$ identity matrix, and similarly for the zero matrices of various sizes. This *normal form* of antisymmetric bilinear forms is an analogue of the Gram-Schmidt theorem that an inner product space has an orthonormal basis, and is proved by an analogous argument.

So if an antisymmetric bilinear form is non-degenerate, then $r \equiv 2n = m$. That is: eq. 3 loses its bottom row and right column consisting of zero matrices, and reduces to the $2n \times 2n$ *symplectic matrix* ω given by

$$(4) \quad \omega := \begin{pmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{1} & \mathbf{0} \end{pmatrix} .$$

Second, the non-degeneracy of ω implies that at any $x \in M$, there is a basis-independent isomorphism ω^\flat from the tangent space T_x to its dual T_x^* . Namely: for any $x \in M$ and $\tau \in T_x$, the value of the 1-form $\omega^\flat(\tau) \in T_x^*$ is defined by

$$(5) \quad \omega^\flat(\tau)(\sigma) := \omega(\sigma, \tau) \quad \forall \sigma \in T_x .$$

This also means that a symplectic structure enables a covector field, i.e. a differential one-form, to determine a vector field. Thus for any function $H : M \rightarrow \mathbb{R}$, so that dH is a differential 1-form on M , the inverse of ω^\flat (which we might write as ω^\sharp), carries dH to a vector field on M , written X_H . This is the key idea whereby in Hamiltonian mechanics, a scalar function H determines a dynamics; cf. Section 2.1.2.

So far, we have noted some implications of ω being non-degenerate. The other part of the definition of a symplectic form (for a manifold), viz. ω being closed, $d\omega = 0$, is also important. We shall see in Section 2.1.3 that it implies that a vector field X on a symplectic manifold M preserves the symplectic form ω (i.e. in more physical jargon: generates (a one-parameter family of) canonical transformations) iff X is Hamiltonian in the sense that there is a scalar function f such that $X = X_f \equiv \omega^\sharp(df)$. Or in terms of the Poisson bracket, with \cdot representing the argument place for a scalar function: $X(\cdot) = X_f(\cdot) \equiv \{\cdot, f\}$.

So much by way of introducing symplectic manifolds. I turn to showing that any cotangent bundle T^*Q is such a manifold. That is: it has, independently of a choice of coordinates or bases, a symplectic structure.

Given a manifold Q ($\dim(Q)=n$) which we think of as the system's configuration space, choose any local coordinate system q on Q , and the natural local coordinates q, p thereby induced on T^*Q . We define the 2-form

$$(6) \quad dp \wedge dq := dp_i \wedge dq^i := \sum_{i=1}^n dp_i \wedge dq^i .$$

In fact, eq. 6 defines the same 2-form, whatever choice we make of the chart q on Q . For $dp \wedge dq$ is the exterior derivative of a 1-form on T^*Q which is defined naturally (i.e. independently of coordinates or bases) from the derivative (also known as: tangent) map of the projection

$$(7) \quad \pi : (q, p) \in T^*Q \mapsto q \in Q.$$

Thus consider a tangent vector τ (not to Q , but) to the cotangent bundle T^*Q at a point $\eta = (q, p) \in T^*Q$, i.e. $q \in Q$ and $p \in T_q^*$. Let us write this as: $\tau \in T_\eta(T^*Q) \equiv T_{(q,p)}(T^*Q)$. The derivative map, $D\pi$ say, of the natural projection π applies to τ :

$$(8) \quad D\pi : \tau \in T_{(q,p)}(T^*Q) \mapsto (D\pi(\tau)) \in T_q .$$

Now define a 1-form θ_H on T^*Q by

$$(9) \quad \theta_H : \tau \in T_{(q,p)}(T^*Q) \mapsto p(D\pi(\tau)) \in \mathbb{R} ;$$

where in this definition of θ_H , p is defined to be the second component of τ 's base-point $(q, p) \in T^*Q$; i.e. $\tau \in T_{(q,p)}(T^*Q)$ and $p \in T_q^*$.

This 1-form is called the *canonical 1-form* on T^*Q . One now checks that in any natural local coordinates q, p , θ_H is given by

$$(10) \quad \theta_H = p_i dq^i .$$

Finally, we define a 2-form by taking the exterior derivative of θ_H :

$$(11) \quad \mathbf{d}(\theta_H) := \mathbf{d}(p_i dq^i) \equiv dp_i \wedge dq^i .$$

One checks that this 2-form is closed (since $\mathbf{d}^2 = 0$) and non-degenerate. So $(T^*Q, \mathbf{d}(\theta_H))$ is a symplectic manifold. Accordingly, $\mathbf{d}(\theta_H)$, or its negative $-\mathbf{d}(\theta_H)$, is called the *canonical symplectic form*, or *canonical 2-form*.

There is a theorem (Darboux's theorem) to the effect that locally, any symplectic manifold "looks like" a cotangent bundle: or in other words, a cotangent bundle is locally a "universal" example of symplectic structure. We will not go into details; but in Section 5.3.4, we will discuss the generalization of this theorem for Poisson manifolds. But first we review, in the next two Subsections, Hamilton's equations, and Noether's theorem.

2.1.2 Geometric formulations of Hamilton's equations

As we already emphasised, the main geometric idea behind Hamilton's equations is that a gradient, i.e. covector, field dH determines a vector field X_H . So to give a geometric formulation of Hamilton's equations at a point $x = (q, p)$ in a cotangent bundle T^*Q , let us write ω^\sharp for the (basis-independent) isomorphism from the cotangent space to the tangent space, $T_x^* \rightarrow T_x$, induced by $\omega := -\mathbf{d}(\theta_H) = dq^i \wedge dp_i$ (cf. eq. 5). Then Hamilton's equations may be written as:

$$(12) \quad \dot{x} = X_H(x) = \omega^\sharp(\mathbf{d}H(x)) = \omega^\sharp(dH(x)) .$$

There are various other formulations. Applying ω^\flat , the inverse isomorphism $T_x \rightarrow T_x^*$, to both sides, we get

$$(13) \quad \omega^\flat X_H(x) = dH(x) .$$

In terms of the symplectic form ω at x , this is: for all vectors $\tau \in T_x$

$$(14) \quad \omega(X_H(x), \tau) = dH(x) \cdot \tau ;$$

or in terms of the contraction (also known as: interior product) $\mathbf{i}_X \alpha$ of a differential form α with a vector field X , with \cdot marking the argument place of $\tau \in T_x$:

$$(15) \quad \mathbf{i}_{X_H} \omega := \omega(X_H(x), \cdot) = dH(x)(\cdot) .$$

More briefly, and now written for any function f , it is:

$$(16) \quad \mathbf{i}_{X_f} \omega = df .$$

Finally, recall the relation between the Poisson bracket and the directional derivative (or the Lie derivative \mathcal{L}) of a function: viz.

$$(17) \quad \mathcal{L}_{X_f} g = dg(X_f) = X_f(g) = \{g, f\} .$$

Combining this with eq. 16, we can state the relation between the symplectic form and Poisson bracket in the form:

$$(18) \quad \{g, f\} = dg(X_f) = \mathbf{i}_{X_f} dg = \mathbf{i}_{X_f}(\mathbf{i}_{X_g} \omega) = \omega(X_g, X_f) .$$

2.1.3 Noether's theorem

The core idea of Noether's theorem, in both the Lagrangian and Hamiltonian frameworks, is that to every continuous symmetry of the system there corresponds a conserved quantity (a first integral, a constant of the motion). The idea of a continuous symmetry is made precise along the following lines: a symmetry is a vector field on the state-space that (i) preserves the Lagrangian (respectively, Hamiltonian) and (ii) "respects" the structure of the state-space.

In the Hamiltonian framework, the heart of the proof is a "one-liner" based on the fact that the Poisson bracket is antisymmetric. Thus for any scalar functions f and H on a symplectic manifold (M, ω) (and so with a Poisson bracket given by eq. 18), we have that at any point $x \in M$

$$(19) \quad X_f(H)(x) \equiv \{H, f\}(x) = 0 \quad \text{iff} \quad 0 = \{f, H\}(x) \equiv X_H(f)(x) .$$

In words: around x , H is constant under the flow of the vector field X_f (i.e. under what the evolution would be if f was the Hamiltonian) iff f is constant under the flow X_H . Thinking of H as the physical Hamiltonian, so that X_H represents the real time-evolution (sometimes called: the dynamical flow), this means: around x , X_f preserves the Hamiltonian iff f is constant under time-evolution, i.e. f is a conserved quantity (a constant of the motion).

But we need to be careful about clause (ii) above: the idea that a vector field “respects” the structure of the state-space. In the Hamiltonian framework, this is made precise as preserving the symplectic form. Thus we define a vector field X on a symplectic manifold (M, ω) to be *symplectic* (also known as: *canonical*) iff the Lie-derivative along X of the symplectic form vanishes, i.e. $\mathcal{L}_X \omega = 0$. (This definition is equivalent to X ’s generating (active) canonical transformations, and to its preserving the Poisson bracket. But I will not go into details about these equivalences: for they belong to the theory of canonical transformations, which, as mentioned, I will not need to develop.)

We also define a *Hamilton system* to be a triple (M, ω, H) where (M, ω) is a symplectic manifold and $H : M \rightarrow \mathbb{R}$, i.e. $M \in \mathcal{F}(M)$. And then we define a (continuous) *symmetry* of a Hamiltonian system to be a vector field X on M that:

- (i) preserves the Hamiltonian function, $\mathcal{L}_X H = 0$; and
- (ii) preserves the symplectic form, $\mathcal{L}_X \omega = 0$.

These definitions mean that to prove Noether’s theorem from eq. 19, it will suffice to prove that a vector field X is symplectic iff it is locally of the form X_f . Such a vector field is called *locally Hamiltonian*. (And a vector field is called *Hamiltonian* if there is a global scalar $f : M \rightarrow \mathbb{R}$ such that $X = X_f$.) In fact, two results from the theory of differential forms, the Poincaré Lemma and Cartan’s magic formula, make it easy to prove this; (for a vector field on any symplectic manifold (M, ω) , i.e. (M, ω) does not need to be a cotangent bundle).

Again writing \mathbf{d} for the exterior derivative, we recall that a k -form α is called:

- (i) *exact* if there is a $(k - 1)$ -form β such that $\alpha = \mathbf{d}\beta$; (cf. the elementary definition of an exact differential);
- (ii) *closed* if $\mathbf{d}\alpha = 0$.

The *Poincaré Lemma* states that every closed form is locally exact. To be precise: for any open set U of M , we define the vector space $\Omega^k(U)$ of k -form fields on U . Then the Poincaré Lemma states that if $\alpha \in \Omega^k(M)$ is closed, then at every $x \in M$ there is a neighbourhood U such that $\alpha|_U \in \Omega^k(U)$ is exact.

Cartan’s magic formula is a useful formula (proved by straightforward calculation) relating the Lie derivative, contraction and the exterior derivative. It says that if X is a vector field and α a k -form on a manifold M , then the Lie derivative of α with respect to X (i.e. along the flow of X) is

$$(20) \quad \mathcal{L}_X \alpha = \mathbf{d}i_X \alpha + i_X \mathbf{d} \alpha \quad .$$

We now argue as follows. Since ω is closed, i.e. $\mathbf{d}\omega = 0$, Cartan's magic formula, eq. 20, applied to ω becomes

$$(21) \quad \mathcal{L}_X \omega \equiv \mathbf{d}i_X \omega + i_X \mathbf{d}\omega = \mathbf{d}i_X \omega \quad .$$

So for X to be symplectic is for $i_X \omega$ to be closed. But by the Poincaré Lemma, if $i_X \omega$ closed, it is locally exact. That is: there locally exists a scalar function $f : M \rightarrow \mathbb{R}$ such that

$$(22) \quad i_X \omega = df \quad \text{i.e.} \quad X = X_f \quad .$$

So for X to be symplectic is equivalent to X being locally Hamiltonian.

Thus we have

Noether's theorem for a Hamilton system If X is a symmetry of a Hamiltonian system (M, ω, H) , then locally $X = X_f$; so by the anti-symmetry of the Poisson bracket, eq. 19, f is a constant of the motion. And conversely: if $f : M \rightarrow \mathbb{R}$ is a constant of the motion, then X_f is a symmetry.

We will see in Section 6.2 that most of this approach to Noether's theorem, in particular the "one-liner" appeal to the anti-symmetry of the Poisson bracket, eq. 19, carries over to the more general framework of Poisson manifolds. For the moment, we mention an example (which we will also return to).

For most Hamiltonian systems in euclidean space \mathbb{R}^3 , spatial translations and rotations are (continuous) symmetries. Let us consider in particular a system we will discuss in more detail in Section 2.3: N point-particles interacting by Newtonian gravity. The Hamiltonian is a sum of two terms, which are each individually invariant under translations and rotations:

- (i) a kinetic energy term K ; though I will not go into details, it is in fact defined by the euclidean metric of \mathbb{R}^3 , and is thereby invariant; and
- (ii) a potential energy term V ; it depends only on the particles' relative distances, and is thereby invariant.

The corresponding conserved quantities are the total linear and angular momentum.⁶

2.2 The road ahead

In this Subsection, four comments will expand on the introductory comment (iv) of Section 1.2, and also give some information about the history of symplectic reduction and about some crucial examples.

⁶By the way, this Hamiltonian is *not* invariant under boosts. But as I said in (iii) of Section 1.2, I restrict myself to time-independent transformations; the treatment of symmetries that "represent the relativity of motion" needs separate discussion.

(1): *Generalizing from Noether's theorem; Poisson manifolds:*

Noether's theorem tells us that a continuous symmetry, i.e. a one-parameter group of symmetries, determines a first integral (i.e. a constant of the motion). So a larger group of symmetries, i.e. a group with several parameters, implies several first integrals. The phase flow is therefore confined to the intersection of the level surfaces of these integrals: an intersection which is in general a manifold. In other words: the simultaneous level manifold of these integrals is an invariant manifold of the phase flow.

It turns out that, in many useful cases, this manifold is also invariant under an appropriately chosen subgroup of the group of symmetries; and that the quotient space, i.e. the set of orbits under the action of this subgroup, is a manifold with a natural structure induced by the original Hamiltonian system that is sufficient to do mechanics in Hamiltonian style. The quotient space is therefore called the 'reduced phase space'.

But in some cases, this natural structure is not a symplectic form, but a (mild) generalization in which the form is allowed to be degenerate; i.e. like eq. 3 rather than eq. 4. A manifold equipped with such a structure need not be a quotient manifold. It can instead be defined in terms of a generalization of the usual Poisson bracket, as defined in terms of the symplectic form by eq. 18.

The key idea is to postulate a bracket, acting on the scalar functions $F : M \rightarrow \mathbb{R}$ on *any* manifold M , and possessing four properties enjoyed by the usual Poisson bracket. One of the properties is anti-symmetry, emphasised in Section 2.1.3's proof of Noether's theorem. The other three are that the postulated bracket, again written $\{, \}$, is: to be bilinear; to obey the Jacobi identity for any real functions F, G, H on M , i.e.

$$(23) \quad \{\{F, H\}, G\} + \{\{G, F\}, H\} + \{\{H, G\}, F\} = 0 \quad ;$$

and to obey Leibniz' rule for products, i.e.

$$(24) \quad \{F, H \cdot G\} = \{F, H\} \cdot G + H \cdot \{F, G\} \quad .$$

We will see in Section 5 that such a bracket, again called 'Poisson bracket', provides a sufficient framework for mechanics in Hamiltonian style. In particular, it induces an anti-symmetric bilinear form that may be degenerate, as in eq. 3. A manifold M equipped with such a bracket is called a *Poisson manifold*.

The allowance of degeneracy means that a Poisson manifold can have *odd* dimension; while we saw in Section 2.1.1 that any symplectic manifold is even-dimensional. On the other hand, this generalized Hamiltonian mechanics will have clear connections with the usual formulation of Section 2.1. The main connection will be the result that any Poisson manifold M is a disjoint union of even-dimensional manifolds, on which M 's degenerate antisymmetric bilinear form restricts to be non-degenerate.⁷

⁷Because of these clear connections, it is natural to still call the more general framework 'Hamiltonian'; as is usually done. But of course this is just a verbal matter.

(2): *Historical roots:*

The theory of symplectic reduction has deep historical roots in the work of classical mechanics' monumental figures. In part, this is no surprise. As mentioned in (i) of Section 1.2, cyclic coordinates underpin the role of symmetry in mechanics, and in particular Noether's theorem. And Newton's solution of the Kepler problem provides an example: witness textbooks' expositions of the transition to centre-of-mass coordinates, and of polar coordinates with the angle being cyclic (yielding angular momentum as the conserved quantity). So it is unsurprising that various results and ideas of symplectic reduction can be seen in the work of such masters as Euler, Lagrange, Hamilton, Jacobi, Lie and Poincaré; for example (as we will see), in Euler's theory of the rigid body.

But the history also holds a surprise. It turns out that Lie's epoch-making work on Lie groups already contained a detailed development of much of the general, modern theory.⁸ The sad irony is that most of Lie's insights were not taken up — and were then repeatedly re-discovered. So this is yet another example (of which the history of mathematics has so many!) of the saying that he who does not learn from history is doomed to repeat it. The consolation is of course that it is often easier, and more fun, to re-discover something than to learn it from another...

Thus it was only from the mid-1960s that the theory, in essentially the form Lie had, was recovered and cast in the geometric language adopted by modern mechanics; namely, by contemporary masters such as Arnold, Kostant, Marsden, Meyer, Smale, Souriau and Weinstein; (cf. this Chapter's first motto). Happily, several of these modern authors are scholars of the history, and even their textbooks give some historical details: cf. [Marsden and Ratiu, 1999, pp. 336–8; 369–370; 430–432], and the notes to each Chapter of Olver [2000, especially p. 427–428]. (Hawkins [2000] is a full history of Lie groups from 1869 to 1926; for Lie, cf. especially its Sections 1.3, 2.5 and Chapter 3, especially 3.2.)

In any case, setting history aside: symplectic reduction has continued since the 1970s to be an active research area in contemporary mechanics, and allied fields such as symplectic geometry. So it has now taken its rightful place as a major part of the contemporary renaissance of classical mechanics: as shown by ...

(3): *Two examples: the rigid body and the ideal fluid:*

Two examples illustrate vividly how symplectic reduction can give new physical understanding, even of time-honoured examples: the rigid body and the ideal fluid — as attested by this Chapter's mottoes. (Section 2.3 will develop a third example, more closely related to philosophy.)

As to the rigid body: we will see (especially in Section 5) that symplectic reduction considerably clarifies the elementary theory of the rigid body (Euler's equations, Euler angles etc.): which, notoriously, can be all too confusing! For simplicity, I shall take the rigid body to be pivoted, so as to set aside translational

⁸The main source is his [1890]. Besides, Arnold [1989, 456] reports that the prototype example of a Poisson manifold, viz. the dual of a finite-dimensional Lie algebra, was already understood by Jacobi.

motion. This will mean that the group of symmetries defining the quotienting procedure will be the rotation group. It will also mean that the rigid body's configuration space is given by the rotation group, since any configuration can be labelled by the rotation that obtains it from some reference-configuration. So in this application of symplectic reduction, the symmetry group (viz. the rotation group) will act on *itself* as the configuration space. This example will also give us our prototype example of a Poisson manifold.

As to the ideal fluid, i.e. a fluid that is incompressible and inviscid (with zero viscosity): this is of course an infinite-dimensional system, and so (as I announced in Section 1.2) outside the scope of this Chapter. So I will not go into any details, but just report the main idea.

The equations of motion of an ideal fluid, Euler's equations, are usually derived either by applying Newton's second law $\mathbf{F} = m\mathbf{a}$ to a small fluid element; or by a heuristic use of the Lagrangian or Hamiltonian approach (as in heuristic classical field theories). But in the mid-1960s, Arnold showed how the latter derivations could be understood in terms of a striking, even beautiful, analogy with the above treatment of the rigid body. Namely, the analogy shows that the configuration space of the fluid is an infinite-dimensional group; as follows. The configuration of an ideal fluid confined to some container occupying a volume $V \subset \mathbb{R}^3$ is an assignment to each spatial position $x \in V$ of an infinitesimal fluid element. Given such an assignment as a reference-configuration, any other configuration can be labelled by the volume-preserving diffeomorphism d from V to V that carries the reference-configuration to the given one, by dragging each fluid element along by d . So given a choice of reference-configuration, the fluid's configuration space is given by the infinite-dimensional group \mathcal{D} of diffeomorphisms $d : V \rightarrow V$: just as the rotation group is the configuration space of a (pivoted) rigid body. \mathcal{D} then forms the basis for rigorous Lagrangian and Hamiltonian theories of an ideal fluid.

These theories turn out to have considerable analogies with the Lagrangian and Hamiltonian theories of the rigid body, thanks to the fact that in both cases the symmetry group forms the configuration space. In particular, Euler's equations for ideal fluids are the analogues of Euler's equations for a rigid body. Besides, these rigorous theories of fluids (and symplectic reduction applied to them) are scientifically important: they have yielded various general theorems, and solved previously intractable problems. (For more details, cf. Abraham and Marsden (1978: Sections 4.4 and 4.6 for the rigid body, and 5.5.8 for the ideal fluid), Arnold (1989: Appendix 2:C to 2:F for the rigid body, and 2:G to 2:L for the ideal fluid), and Marsden and Ratiu (1999: Chapters 1.4 and 15 for the rigid body, and 1.5, p. 266, for the ideal fluid).)

(4): *Philosophical importance:*

Symplectic reduction is also, I submit, philosophically important; in at least two ways. The first way is specific: it illustrates some methodological morals about how classical mechanics analyses problems. I develop this theme in (Butterfield 2005). The second way is more general: the theory, or rather various applications

of it, is directly relevant to disputes in the philosophy of space and time, and of mechanics. This relevance is recognized in contemporary philosophy of physics. So far as I know, the authors who develop these connections in most detail are Belot and Earman. They discuss symplectic reduction in connection with such topics as:

- (i) the treatment of symmetries, including gauge symmetries;
- (ii) the dispute between absolute and relationist conceptions of space and time; and
- (iii) the interpretation of classical general relativity (a topic which connects (i) and (ii), and bears on heuristics for quantum gravity).

Thus Belot [1999; 2000; 2001; 2003; 2003a] and Earman [2003] discuss mainly (i) and-or (ii); Belot and Earman (2001) discusses (iii). For (i) and (ii), I also recommend Wallace [2003].

But these papers have a demanding pre-requisite: they invoke, but do not expound, the theory of symplectic reduction. They also discuss infinite-dimensional systems (especially classical electromagnetism and general relativity), without developing finite-dimensional examples like the rigid body. Indeed, there is, so far as I know, no article-length exposition of the theory which is not unduly forbidding for philosophers. So I aim to give such an exposition, to help readers of papers such as those cited.⁹

As an appetizer for this exposition, I will first (in Section 2.3) follow Belot in presenting the general features of a finite-dimensional symplectic reduction which has vivid philosophical connections, viz. to the absolute vs. relationist debate. This example concerns a system of point-particles in Euclidean space, either moving freely or interacting by a force such as Newtonian gravity. (The symmetries defining the quotienting procedure are given by the Euclidean group of translations and rotations.) For philosophers, this will be a good appetizer for symplectic reduction, since it sheds considerable light on relationism about space of the sort advocated by Leibniz and Mach.

2.3 *Appetizer: Belot on relationist mechanics*

2.3.1 *Comparing two quotienting procedures*

In several papers, Belot discusses how symplectic reduction bears on the absolute-vs.-relational debate about space. I shall pick out one main theme of his discussions: the comparison of a relational classical mechanical theory with what one

⁹As I said in Section 1.2, my material is drawn from the books by Abraham and Marsden, Arnold, Marsden and Ratiu, and Olver. More precisely, I will mostly draw on: Abraham and Marsden [1978, Sections 3.1–3.3, 4.1–4.3], Arnold [1989: Appendices 2, 5 and 14], Marsden and Ratiu [1999, Chapters 9–13] and Olver [2000, Chapter 6]. And much of what follows — in spirit, and even in letter — is already in Lie [1890]! As a (non-philosophical) introduction to symplectic reduction, I also recommend Singer (2001). It is at a yet more elementary level than what follows; e.g. it omits Poisson manifolds and co-adjoint representations.

gets by quotienting the orthodox absolutist (also called a ‘substantivalist’) classical mechanics, by an appropriate symmetry group. His main contention — which I endorse — is that this comparison sheds considerable light on relationism: on both its motivation, and its advantages and disadvantages.¹⁰

Belot’s overall idea is as follows. Where the relationist admits one possible configuration, as (roughly) a specification of all the distances (and thereby angles) between all the parts of matter, the absolutist (or substantivalist) sees an infinity of possibilities: one for each way the relationist’s configuration (a *relative configuration*) can be embedded in the absolute space. This makes it natural to take the relationist to be envisaging a mechanics which is some sort of “quotient” of the absolutist’s mechanics.

In particular, on the traditional conception of space as Euclidean (modelled by \mathbb{R}^3), each of the relationist’s relative configurations corresponds to an equivalence class of absolutist configurations (i.e. embeddings of arrangements of matter into \mathbb{R}^3), with the members of the class related by spatial translations and rotations, i.e. elements of the Euclidean group. In the jargon of group actions, to be developed in Section 4: the Euclidean group acts on the set of all absolutist configurations, and a relative configuration corresponds to an orbit of this action. So it is natural to take the relationist to be envisaging a mechanics which quotients this action of the Euclidean group, to get a *relative configuration space*. A relationist mechanics, of Lagrangian or Hamiltonian type, is then to be built up on this space of relative configurations.

But as Belot emphasises, one can instead consider quotienting the absolutist’s state-space — i.e. in a Hamiltonian framework, the phase space — rather than their configuration space. Indeed, this is exactly what one does in symplectic reduction. In particular, the Euclidean group’s action on the absolutist’s configuration space, Q say, can be lifted to give an action on the cotangent bundle T^*Q ; which is accordingly called the ‘cotangent lift’. One can then take the quotient, i.e. consider the orbits into which T^*Q is partitioned by the cotangent lift.

We thus have two kinds of theories to compare: (i) the relationist theories, built up from the relative configuration space; which for the sake of comparison with symplectic reduction we take to be Hamiltonian, rather than Lagrangian; (ii) theories obtained by quotienting “later”, i.e. quotienting the absolutist’s cotangent bundle.

I will now spell out this comparison. But I will not try to summarize Belot’s more detailed conclusions, about what such a comparison reveals about the advantages and disadvantages of relationism. They are admirably subtle, and so defy summary: they can mainly be found at his [2000, p. 573–574, 582; 2001, Sections VIII to X]. (Rovelli (this volume) also discusses relationism.)

As befits an appetizer, I will also (like Belot) concentrate on as simple a case as possible: a mechanics of N point-particles, which is to assume a Euclidean

¹⁰The main references are Belot [1999; 2001; 2003, Sections 3.5, 5]. Cf. also his [2000, Sections 4 to 5.3; 2003a, Section 6]. Though I recommend all these papers, the closest template for what follows is [2001, Section VI et seq.].

spatial geometry. Of course, the absolutist make this assumption by postulating a Euclidean space; but for the relationist, the assumption is encoded in constraints relating the various inter-particle distances. The main current example of a relationist mechanics of such a system is due to Barbour and Bertotti (1982), though they develop it in the Lagrangian framework; (to be precise, in terms of Jacobi's principle). Belot also discusses other relational theories, including field theories, i.e. theories of infinite systems; some of them also due to Barbour, and in a Lagrangian framework. But in this Section I only consider N point-particles.

Also, I will also not discuss boosts, though of course the relationist traditionally proposes to identify any two absolutist states of motion related by a boost. In terms of group actions, this means I will consider quotienting by an action of the euclidean group, but not the Galilei group. (Cf. how I set aside time-dependent transformations already in (iii) of Section 1.2.) I will also postpone to later Sections technical details, even when our previous discussion makes them accessible.

Finally, a warning to avoid later disappointment! The later Sections will not include a full analysis of the euclidean group's actions on configuration space and phase space, and their quotients. That would involve technicalities going beyond an appetizer. Instead (as mentioned at the end of Section 1.2), the material in later Sections is chosen so as to lead up to Section 7's theorem, the Lie-Poisson reduction theorem, about quotienting the phase space of a system whose configuration space is a Lie group. Further reasons for presenting the material for this theorem will be given in Section 5.1.

2.3.2 The spaces and group actions introduced

Let us begin by formulating the orthodox absolutist mechanics of N point-particles interacting by Newtonian gravity, together with the action of the Euclidean group.

Each point-particle occupies a point of \mathbb{R}^3 , so that the configuration space Q is \mathbb{R}^{3N} : $\dim(Q) = 3N$. So the phase space for Hamiltonian mechanics will be the cotangent bundle $T^*Q \ni (q, p)$: $\dim(T^*Q) = 6N$.

The Hamiltonian is a sum of kinetic and potential terms, K and V . K depends only on the ps , and V only on the qs . In cartesian coordinates, with i now labelling particles $i = 1, \dots, N$ rather than degrees of freedom, we have the familiar expressions:

$$(25) \quad H(q, p) = K(p) + V(q) \quad \text{with} \quad K = \sum_i \frac{\mathbf{p}_i^2}{2m_i}, \quad V(q) = G \sum_{i < j} \frac{m_i m_j}{\|\mathbf{q}_i - \mathbf{q}_j\|}$$

where m_i are the masses and G is the gravitational constant.^{11,12}

¹¹From the broader philosophical perspective, the most significant feature of eq. 25 is no doubt the fact that the potential is a sum of all the two-body potential energies for the configuration $q \in Q$: there are no many-body interactions.

¹²Incidental remark. In fact, the kinetic energy can be represented by a metric g on the configuration space. For Hamiltonian mechanics, this means that the kinetic energy scalar K on the cotangent bundle T^*Q can be defined by applying Q 's metric g to the projections of the momenta p , where at each point $(q, p) \in T^*Q$ the projection is made with the preferred

The euclidean group E (aka: $E(3)$) is the group (under composition) of translations, rotations and reflections on \mathbb{R}^3 . But since we will be interested in continuous symmetries, we will ignore reflections, and so consider the subgroup of orientation-preserving translations and rotations; i.e. the component of the group connected to the identity transformation (which I will also write as E). This is a Lie group, i.e. a group which is also a manifold, with the group operations smooth with respect to the manifold structure. Section 3 will give formal details. Here we just note that we need three real numbers to specify a translation ($\mathbf{x} = (x, y, z)$), and three to specify a rotation (two for an axis, and one for the angle through which to rotate); and accordingly, it is unsurprising that as a manifold, the dimension of E is 6: $\dim(E) = 6$.

E acts in the obvious sense on \mathbb{R}^3 . For example, if $g \in E$ is translation by $\mathbf{x} \in \mathbb{R}^3$, g induces the map $\mathbf{q} \in \mathbb{R}^3 \mapsto \mathbf{q} + \mathbf{x}$. Similarly for a rotation induces: again, Section 3 will give a formal definition.

Now let E act in this way on each of the N factor spaces \mathbb{R}^3 of our system's configuration manifold $Q = \mathbb{R}^{3N}$. This defines an action Φ on Q : i.e. for all $g \in E$, there is a map $\Phi_g : Q \rightarrow Q$. For example, for g a translation by $\mathbf{x} \in \mathbb{R}^3$, we have

$$(27) \quad \Phi_g : (\mathbf{q}_j) = (\mathbf{q}_1, \dots, \mathbf{q}_N) \in Q \mapsto (\mathbf{q}_1 + \mathbf{x}, \dots, \mathbf{q}_N + \mathbf{x}) \in Q ;$$

and similarly for rotations. Since the potential function $V : Q \rightarrow \mathbb{R}$ of eq. 25 depends only on inter-particle distances, each map $\Phi_g : Q \rightarrow Q$ is a symmetry of the potential; i.e. we have $V(\Phi_g(q)) = V(q)$.

The action Φ (i.e. the assignment $g \in E \mapsto \Phi_g$) induces an action of E on $T^*Q = T^*\mathbb{R}^{3N}$, called the *cotangent lift* of Φ to T^*Q , and usually written as Φ^* ; so that we have for each $g \in E$ a lifted map $\Phi_g^* : T^*Q \rightarrow T^*Q$. Again, the details can wait till later (Section 4). But the idea is that each map Φ_g on Q is smooth, and so maps curves to curves, and so vectors to vectors, and so covectors to covectors, and so on.

Unsurprisingly, each of the lifted maps $\Phi_g^* : T^*Q \rightarrow T^*Q$ leaves the potential V , now considered as a scalar on T^*Q , invariant: i.e. we have $V(\Phi_g^*(q, p)) = V(q, p) \equiv V(q)$. But furthermore, each of the lifted maps Φ_g^* is a symmetry of the Hamilton system, in our previous sense (Section 2.1.3). That is: Φ_g^* preserves the Hamiltonian (indeed the kinetic and potential terms are separately invariant); and it preserves the symplectic structure. This means the dynamics is invariant under the action of all $g \in G$: the dynamical histories of the system through (q, p) and through $\Phi_g^*(q, p)$ match exactly at each time. They are qualitatively indistinguishable: in contemporary metaphysical jargon, they are *duplicates*.

At this point, of course, we meet the absolute-vs.-relational debate about space. The absolutist asserts, and the relationist denies, that there being two such in-

isomorphism $\omega^\sharp : T_q^* \rightarrow T_q$; (cf. eq. 12). That is: —

$$(26) \quad K : (q, p) \in T^*Q \mapsto g_q(\omega^\sharp(p), \omega^\sharp(p)) .$$

distinguishable possibilities makes sense.¹³ So the relationist, presented with the theory above, says we should cut down the space of possibilities. As I said in Section 2.3.1, it is natural to make this precise in terms of quotienting the action of the euclidean group: a set of absolutist possibilities related one to another by elements of the euclidean group form an equivalence class (an orbit) which is to represent *one* relationist possibility.

But here we need to distinguish two different quotienting procedures. I will call them *Relationism* and *Reductionism* (with capital R's), since the former is close to both traditional and contemporary relationist proposals, and the latter is an example of the orthodox idea of symplectic reduction. As I said in Section 2.3.1, the main difference will be that:

- (i) Relationism performs the quotient on E 's action on the configuration space Q ; the set of orbits form a *relative configuration space*, on which the relationist proposes to build a dynamics, whether Lagrangian or Hamiltonian — yielding in the latter case, a *relative phase space*; whereas
- (ii) Reductionism performs the quotient on E 's action on the usual phase space T^*Q , the set of orbits forming a *reduced phase space*.

Since our discussion adopts the Hamiltonian framework, it will not matter for what follows, that Relationism, as defined, can adopt the Lagrangian framework, while Reductionism is committed to the Hamiltonian one. What *will* matter is that (i) and (ii) make for phase spaces of different dimensions; the reduced phase space has six more dimensions than the relative phase space. The “dimension gap” is six.

We will see that four of the six variables that describe these dimensions are constants of the motion; the other two vary with time. And for certain choices of values of the constants of the motion (roughly: no rotation), the time-varying variables drop out, and the dynamics according to the Reductionist theory simplifies so as to coincide with that of the Relationist theory. In other words: if we impose no rotation, then the heterodox Relationist dynamics matches the conventional Reduced dynamics.

2.3.3 The Relationist procedure

The Relationist seeks a mechanics based on the *relative configuration space* (RCS). An element of the RCS is to be a pattern of inter-particle distances and angles that is geometrically possible, i.e. compatible with the N particles being embedded in

¹³The *locus classicus* for this debate is of course the Leibniz-Clarke correspondence, though the protagonists' argumentation is of course sometimes theological. Clarke the absolutist maintains that there are many possible arrangements of bits of matter in space consistent with a specification of all relative distances, saying ‘if [the mere will of God] could in no case act without a pre-determining cause ... this would tend to take away all power of choosing, and to introduce fatality.’ Leibniz claims there is only one such arrangement: ‘those two states ... would not at all differ from one another. Their difference therefore is only to be found in our chimerical supposition of the reality of space in itself.’

\mathbb{R}^3 . So, roughly speaking, an element of the RCS is a euclidean configuration, modulo isometries; and the RCS will be the set of orbits \mathbb{R}^{3N}/E .

Even before giving a more precise statement, we can state the “punchline” about dimensions, as follows. Since $\dim(E) = 6$, quotienting by E subtracts six dimensions: that is, the dimension of the RCS will be $3N-6$.

But we need to be more precise about the RCS. For the orbits and quotient spaces to be manifolds, and for dimensions to add or subtract in this simple way, we need to excise two classes of “special” points from \mathbb{R}^{3N} , before we quotient. (But I postpone till Section 4 the technical rationale for these excisions.)

Let $\delta_Q \subset \mathbb{R}^{3N}$ be the set of configurations which are symmetric: i.e. each is fixed by some element of E (other than the identity element!). Any configuration in which all the point-particles are collinear provides an example: the configuration is fixed by any rotation about the line as axis. Let Δ_Q be the set of collision configurations; i.e. configurations in which two or more particles are coincident in the usual configuration space \mathbb{R}^{3N} . (The Q subscripts will later serve as a reminder that these sets are sets of configurations.) δ_Q and Δ_Q are both of measure zero in \mathbb{R}^{3N} . Excise both of them, and call the resulting space, which is again of dimension $3N$: $Q := \mathbb{R}^{3N} - (\delta_Q \cup \Delta_Q)$.

δ_Q and Δ_Q are each closed under the action of E . That is, each is a union of orbits: a euclidean transformation of a symmetric (collision) configuration is also symmetric (collision). So E acts on Q . Now quotient Q by E . Q/E is the Relationist’s RCS. Since $\dim(E) = 6$, we have: $\dim(Q/E) = 3N-6$.

These $3N - 6$ variables encode all of a (relative) configuration’s particle-pair relative distances, $r_{ij} \in \mathbb{R}$ (with i, j labelling particles). Note that there are $N(N - 1)/2$ such relative distances; and for $N > 4$, this is greater than $3N - 6$: (for $N \gg 4$, it is much greater). So the relative distances, though physically intuitive, give an *over-complete* set of coordinates on Q/E . (So they cannot be freely chosen: there are *constraints* between them.)

So the Relationist seeks a mechanics that uses this RCS. Newton’s second law being second-order in time means that she will also need quantities like velocities (in a Lagrangian framework) or like momenta (in a Hamiltonian framework). For the former, she will naturally consider the $N(N - 1)/2$ relative velocities $\dot{r}_{ij} := \frac{d}{dt}r_{ij}$; and for the latter, the corresponding momenta $p_{ij} := \frac{\partial L}{\partial \dot{r}_{ij}}$. Again, she must beware of constraints. The tangent and cotangent bundles built on her RCS Q/E will each have dimension $2(3N - 6) = 6N - 12$. So again, for $N > 4$, the number $N(N - 1)/2$ of relative velocities \dot{r}_{ij} , or of relative momenta p_{ij} , is greater than the number of degrees of freedom concerned; and for $N \gg 4$, it is much greater. So again, the relative velocities or relative momenta are over-complete: there are constraints.

On the other hand, if the Relationist uses only these relative quantities, r_{ij} and either \dot{r}_{ij} or p_{ij} (or “equivalent” coordinates on $T(Q/E)$ or $T^*(Q/E)$ that are not over-complete), she faces a traditional problem — whatever the other details of her theory. At least, she faces a problem if she hopes for a deterministic theory which is empirically equivalent to the orthodox absolutist theory. I will follow tradition

and state the problem in terms of relative velocities rather than momenta.

The problem concerns rotation; (and herein lies the strength of Newton's and Clarke's position in the debate against Leibniz). For according to the absolutist theory two systems of point-particles could match with respect to all relative distances and relative velocities, and yet have different future evolutions; so that a theory allowing the same possibilities as the absolutist one, yet using only these relative quantities (or "equivalent" variables) would have to be indeterministic.

The simplest example is an analogue for point-particles of Newton's two globes thought-experiment. Thus the systems could each consist of just two point-particles with zero relative velocity. One system could be non-rotating, so that the point-particles fall towards each other under gravity; while the other system could be rotating about an axis normal to the line between the particles, and rotating at just such a rate as to balance the attractive force of gravity.

The Relationist has traditionally replied that they do *not* hope for a theory empirically equivalent to the absolutist one. Rather, they envisage a mechanics in which, of the two systems mentioned, only the non-rotating evolution is possible: more generally, a mechanics in which the universe as a whole must have zero angular momentum. Originally, in authors like Leibniz and Mach, this reply was a promissory note. But modern Relationist theories such as Barbour and Bertotti's [1982] have made good the promise; and they have been extended well beyond point-particles interacting by Newtonian gravity. Besides, since the universe seems in fact to be non-rotating, these theories can even claim to be empirically adequate, at least as regards this principal difference from absolutist theories.¹⁴

But it is not my brief to go into these theories' details, except by way of comparison with a quotiented version of the absolutist theory: cf. Section 2.3.4.

2.3.4 *The Reductionist procedure*

The Reductionist's main idea is to quotient only after passing to the orthodox phase space for N point-particles, i.e. the cotangent bundle $T^*\mathbb{R}^{3N}$ of \mathbb{R}^{3N} . So the idea is to consider $(T^*\mathbb{R}^{3N})/E$, i.e. the quotient of $T^*\mathbb{R}^{3N}$ by the cotangent-lifted action Φ^* of the euclidean group E .

More precisely, we again proceed by first excising special points that would give technical trouble. But now the points to be excised are in the cotangent bundle $T^*\mathbb{R}^{3N}$, not in \mathbb{R}^{3N} . So let $\delta \subset T^*\mathbb{R}^{3N}$ be the set of phase space states whose configurations are symmetric (in the sense of Section 2.3.3's δ_Q). Let $\Delta \subset T^*\mathbb{R}^{3N}$ be the set of collision points; i.e. states in which two or more particles are coincident in the configuration space \mathbb{R}^{3N} . Both δ and Δ are of measure zero. Excise both of them, and call the resulting phase space, which is again of dimension $6N$: $M := T^*\mathbb{R}^{3N} - (\delta \cup \Delta)$.

¹⁴An advocate of the absolutist theory might say that it is odd to make what seems a contingent feature of the universe, non-rotation, a principle of mechanics; and the Relationist might reply that their view has the merit of predicting that the universe does not rotate! I fear there are no clear criteria for settling this methodological dispute; anyway, I will not pursue it.

δ and Δ are each closed under the cotangent-lifted action of E on $T^*\mathbb{R}^{3N}$. That is, each is a union of orbits: the cotangent lift of a euclidean transformation acting on a phase space state with a symmetric (collision) configuration yields a state which also has a symmetric (collision) configuration. So E acts on M . Now quotient M by E , getting $\bar{M} := M/E$. This is called *reduced phase space*. We have: $\dim(\bar{M}) = \dim(M) - \dim(E) = 6N - 6$.

As emphasised at the end of Section 2.3.2, \bar{M} has six more dimensions than the corresponding Relationist phase space (whether the velocity phase space (tangent bundle) or the momentum phase space (cotangent bundle)). The dimension of *those* phase spaces is $2(3N - 6) = 6N - 12$. Indeed, we can better understand both the reduced phase space \bar{M} and Relationist phase spaces by considering this “dimension gap”. There are two extended comments to make.

(1): *Obtaining the Relationist phase space:*

We can obtain the Relationist momentum phase space from our original phase space M . Thus let M_0 be the subspace of M in which the system has total linear momentum and total angular momentum both equal to zero. Since these are constants of the motion, M_0 is dynamically closed and so supports a Hamiltonian dynamics given just by restriction of the original dynamics. With linear and angular momentum each contributing three real numbers, $\dim(M_0) = \dim(M) - 6 = 6N - 6$. Furthermore, M_0 is closed under (is a union of orbits under) the cotangent-lifted action of E . So let us quotient M_0 by this action of E , and write $\bar{M}_0 := M_0/E$. Then $\dim(\bar{M}_0) = 6N - 6 - 6 = 6N - 12$.

Now recall that this is the dimension of the phase space of the envisaged Relationist theory built on the RCS Q/E . And indeed, as one would hope: \bar{M}_0 is the Hamiltonian version of Barbour and Bertotti’s 1982 Relational theory; (recall that they work in a Lagrangian framework).

That is: \bar{M}_0 is a symplectic manifold, and points in \bar{M}_0 are parametrized by all the particle-to-particle relative distances and relative velocities. There is a deterministic dynamics which matches that of the original absolutist theory, once the original dynamics is projected down to Section 2.3.3’s relative configuration space Q/E .

In short: the vanishing total linear and angular momenta mean that an initial state comprising only relative quantities *is* sufficient to determine all future relative quantities.

(2): *Decomposing the Reductionist reduced phase space:*

Let us return to the reduced phase space \bar{M} . The first point to make is that since the Hamiltonian H on M , or indeed on $T^*\mathbb{R}^{3N}$, is invariant under the cotangent-lifted action of E , the usual dynamics on M projects down to $\bar{M} = M/E$. That is: the reduced phase space dynamics captures all the E -invariant features of the usual dynamics.

In fact, \bar{M} is a *Poisson manifold*. So it is our first example of the more general framework for Hamiltonian mechanics announced in (1) of Section 2.2. Again, I postpone technical detail till later (especially Sections 5.1 and 5.2.4). But the idea

is that a Poisson manifold has a *degenerate* antisymmetric bilinear map, which implies that the manifold is a disjoint union of symplectic manifolds. Each symplectic manifold is called a *leaf* of the Poisson manifold. The leaves' symplectic structures “mesh” with one another; and within each leaf there is a conventional Hamiltonian dynamics.

Even without a precise definition of a Poisson manifold, we can describe how M is decomposed into symplectic manifolds, each with a Hamiltonian dynamics. Recall that we have: $\dim(\bar{M}) = \dim(M) - \dim(E) = 6N - 6$. This breaks down as:

$$(28) \quad 6N - 6 = (6N - 12) + 3 + 3 = 2(3N - 6) + 3 + 3 =: \alpha + \beta + \gamma$$

where the right hand side defines α, β, γ respectively as $2(3N - 6), 3$ and 3 . In terms of \bar{M} , this means the following.

- (i) α corresponds to (1)'s \bar{M}_0 , i.e. to $T^*(Q/E)$. As discussed, $3N - 6$ variables encode all the particle-pair relative distances; and the other $3N - 6$ variables encode all the particle-pair relative momenta.

The six extra variables additional to these $6N-12$ relative quantities consist of: four constants of the motion, and two other variables which are dynamical, i.e. change in time.

- (ii) β stands for three of the four constants of the motion: viz. the three variables that encode the total linear momentum of the system, i.e. the momentum of the centre of mass. These constants of the motion are “just parameters” in the sense that: (a) not only does specifying a value for all three of them fix a surface, i.e. a $(6N - 9)$ -dimensional hypersurface in \bar{M} , on which there is a Hamiltonian dynamics; also (b) this Hamiltonian and symplectic structure is independent of the values we specify.¹⁵
- (iii) γ stands for the three variables that encode the total angular momentum of the system. One of these is a fourth constant of the motion, viz. the magnitude L of the total angular momentum. The other two time-varying quantities fix a point on a sphere (2-sphere) of radius L , encoding the direction of the angular momentum of the system *in a frame rotating with it*. The situation is as in the elementary theory of the rigid body: though the total angular momentum relative to coordinates fixed in space is a constant of the motion (three constant real numbers), the total angular momentum relative to the body is constant only in magnitude (one real number L), not in direction. This will be clearer in Section 5 onwards, when we describe the Poisson manifold structure in the theory of the rigid body. For the moment,

¹⁵As mentioned at the end of Section 2.3.1, the relationist traditionally proposes to identify absolutist states of motion that differ just by the value of the total momentum. And indeed, the proposal can be implemented by considering an action of the Galilean group on the absolutist phase space M , and identifying points related by Galilean boosts. For discussion and references, cf. Belot [2000, Section 5.3].

there are two main comments, (a) and (b), to make about the N particle system:

- (a) If we specify L , in addition to the momentum of the centre of mass of the system, we get a $(6N-10)$ -dimensional hypersurface in \bar{M} , on which (as in (ii)) there is a Hamiltonian dynamics. So we can think of \bar{M} as consisting of the four real-parameter family of these hypersurfaces, with each point of each hypersurface being equipped with a sphere of radius L ; (subject to a qualification in (b) below).

Note that here ‘each point being equipped’ does not mean that the sphere gives the extra dimensions that would constitute \bar{M} as a fibre bundle; (there would be two dimensions lacking). Rather: in the point’s representation by $6N-10$ real numbers, two of the numbers can be taken to represent a point on a sphere.

- (b) But unlike the situation for β in (ii) above, the Hamiltonian dynamics on such a hypersurface *depends* on the value of L . In particular, if $L=0$ the sphere representing the body angular momentum is degenerate: it is of radius zero, and the other two time-varying quantities drop out. A point in the hypersurface is represented by $6N-12$ real numbers; i.e. the hypersurface is $6N-12$ -dimensional.

Now recall from Section 2.3.3 or (1) above that $6N-12$ is the dimension of the phase space of the envisaged Relationist theory built on the RCS Q/E . And indeed, just as one would hope: the hypersurface with $L=0$ and also with vanishing linear momentum, with its dynamics, *is* the symplectic manifold and dynamics that is the Hamiltonian version of Barbour and Bertotti’s 1982 Relational theory of N point-particles. In terms of (1)’s notation, this hypersurface is \bar{M}_0 .

We can sum up this comparison as follows. On this hypersurface \bar{M}_0 , the dynamics in the reduced phase space coincides with the dynamics one obtains for the relative variables, if one arbitrarily embeds their initial values in the usual absolutist phase space $T^*\mathbb{R}^{3N}$, subject to the constraint that the total angular and linear momenta vanish, and then reads off (just by projection) their evolution from the usual evolution in $T^*\mathbb{R}^{3N}$.

2.3.5 Comparing the Relationist and Reductionist procedures

In comparing the Relationist and Reductionist procedures, I shall just make just two extended comments, and refer to Belot for further discussion. The gist of both comments is that Reductionism suffices: Relationism is not needed. The first is a commonplace point; the second is due to Belot.

2.3.5.1 Reductionism allows for rotation The first comment reiterates the Reductionist’s ability, and the Relationist’s inability, to endorse Newton’s globes (or bucket) thought-experiment. The Reductionist can work in either

- (i) the $(6N - 6)$ -dimensional phase space $\bar{M} = M/E$; or
- (ii) the $(6N - 9)$ -dimensional hypersurface got from (i) by specifying the centre of mass' linear momentum; or
- (iii) the $(6N - 10)$ -dimensional hypersurface got from (ii) by also specifying a non-zero value of L .

In all three cases, the Reductionist can describe rotation in a way that the Relationist with their $(6N - 12)$ -dimensional space cannot. For she has to hand the three extra non-relative variables (L and two others) that describe the rotation of the system as a whole. (Incidentally: that they describe the system as a whole is suggested by there being just three of them, whatever the value of N .) In particular, she can distinguish states of rotation and non-rotation ($L = 0$), in the sense of endorsing the distinctions advocated by the globes and bucket thought-experiments.

The Reductionist can also satisfy a traditional motivation for relationism, which concerns general philosophy, rather than the theory of motion. It is especially associated with Leibniz: namely, our theory (or our metaphysics) should not admit distinct but utterly indiscernible possibilities. One might well ask why we should endorse this “principle of the identity of indiscernibles” for possibilities rather than objects. For Leibniz himself, the answer lies (as Belot’s [2001] brings out) in his principle of sufficient reason, and ultimately in theology.

But in any case the Reductionist can satisfy the requirement. Agreed, the usual absolutist theory, cast in $T^*\mathbb{R}^{3N}$ (or if you prefer, $M = T^*\mathbb{R}^{3N} - (\delta \cup \Delta)$) has nine variables that describe (i) the position of the centre of mass, (ii) the orientation of the system about its centre of mass, and (iii) the system’s total linear momentum: i.e. three variables, a vector in \mathbb{R}^3 , for each of (i)-(iii). So the usual absolutist theory has a nine-dimensional “profligacy” of distinct but indiscernible possibilities. But as we have seen, the Reductionist quotients by the action of the euclidean group E , and so works in $\bar{M} = M/E$: which removes the profligacy about (i) and (ii).

As to (iii), I agree that for all I have said, a job remains to be done. The foliation of \bar{M} by a three real-parameter family of $(6N - 9)$ -dimensional hypersurfaces, labelled by the system’s total linear momentum, codifies the profligacy — but does not eliminate it. But as I mentioned above (cf. footnote 15), the Reductionist can in fact quotient further, by considering the action of Galilean boosts and identifying phase space points that differ by a boost; i.e. defining orbits transverse to these hypersurfaces.

2.3.5.2 Analogous reductions in other theories I close my philosophers’ appetizer for symplectic reduction by summarizing some general remarks of Belot’s [2001, Sections VIII–IX]; cf. also his [2003a, Sections 12, 13]. They are about how our discussion of relational mechanics is typical of many cases; and how symplectic reduction can be physically important. I label them (1)–(3).

(1): *A general contrast: when to quotient:*

The example of N point-particles interacting by Newtonian gravity is typical of a large class of cases (infinite-dimensional, as well as finite-dimensional). There is a configuration space Q , acted on by a continuous group G of symmetries, which lifts to the cotangent bundle T^*Q , with the cotangent lift leaving invariant the Hamiltonian, and so the dynamics. So we can quotient T^*Q by G to give a reduced theory. (There is a Lagrangian analogue; but as above, we set it aside.) But there is also some motivation for quotienting G 's action on Q , irrespective of how we then go on to construct dynamics. Let us adopt 'relationism' as a mnemonic label for whatever motivates quotienting the configuration space. Then with suitable technical conditions assumed (recall our excision of δ and Δ), we will have:

(i) for the reduced Hamiltonian theory: $\dim((T^*Q)/G) = 2 \dim Q - \dim G$;

(ii) for the relationist theory, in a Lagrangian or Hamiltonian framework:
 $\dim(T(Q/G)) = \dim(T^*(Q/G)) = 2(\dim Q - \dim G)$

So we have in the reduced theory, $\dim G$ variables that do not occur in the relationist theory: let us call them 'non-relational variables'.

(2): *The non-relational variables:*

Typically, these non-relational variables represent global, i.e. collective, properties of the system. That is unsurprising since the number, $\dim G$, of these variables is independent of the number of degrees of freedom of the system ($\dim Q$, or $2\dim Q$ if you count rate of change degrees of freedom separately).

Some of these variables are conserved quantities, which arise (by Noether's theorem) from the symmetries. Furthermore, there can be specific values of the conserved quantities, like the vanishing angular momentum of Section 2.3.4, for which the reduced theory collapses into the relationist theory. That is, not only are the relevant state spaces of equal dimension; but also their dynamics agree.

(3): *The reduced theory:*

Typically, the topology and geometry of the reduced phase space $(T^*Q)/G$, and the Hamiltonian function on it, $\bar{H} : (T^*Q)/G \rightarrow \mathbb{R}$ say, are more complex than the corresponding features of the unreduced theory on T^*Q . In particular, the reduced Hamiltonian \bar{H} typically has potential energy terms corresponding to forces that are absent from the unreduced theory. But this should not be taken as necessarily a defect, for two reasons.

First, there are famous cases in which the reduced theory has a distinctive motivation. One example is Hertz' programme in mechanics, viz. to "explain away" the apparent forces of our macroscopic experience (e.g. gravity) as arising from reduction of a theory that has suitable symmetries. (The programme envisaged cyclic variables for microscopic degrees of freedom that were unknown to us; cf. Lutzen [1995; 2005].) Another famous example is the Kaluza-Klein treatment of the force exerted on a charged particle by the electromagnetic field. That is: the

familiar Lorentz force-law describing a charged particle’s motion in four space-time dimensions can be shown to arise by symplectic reduction from a theory postulating a spacetime with a fifth (tiny and closed) spatial dimension, in which the particle undergoes straight-line motion. Remarkably, the relevant conserved quantity, viz. momentum along the fifth dimension, can be identified with electric charge; so that the theory can claim to explain the conservation of electric charge. (This example generalizes to other fields: for details and references, cf. Marsden and Ratiu [1999, Section 7.6].)

Second, the reduced theory need not be so complicated as to be impossible to work with. Indeed, these two examples prove this point, since in them the reduced theory is entirely tractable: for it is the familiar theory — that one might resist abandoning for the sake of the postulated unreduced theory.¹⁶ Besides, Belot describes how, even when the reduced theory seems complicated (and not just because it is unfamiliar!), the general theory of symplectic reduction, as developed over the last forty years, has shown that one can often “do physics” in the reduced phase space: and that, as in the Kaluza-Klein example, the physics in the reduced phase space can be heuristically, as well as interpretatively, valuable.

3 SOME GEOMETRIC TOOLS

So much by way of an appetizer. The rest of the Chapter, comprising this Section and the next four, is the five-course banquet! This Section expounds some modern differential geometry, especially about Lie algebras and Lie groups. Section 4 takes up actions by Lie groups. Then Section 5 describes Poisson manifolds as a generalized framework for Hamiltonian mechanics. As I mentioned in (2) of Section 2.2, Lie himself developed this framework; so in effect, he knew everything in these two Sections —so it is a true (though painful!) pun to say that these three Sections give us the “Lie of the land”. In any case, these two Sections will prepare us for Section 6’s description of symmetry and conservation in terms of momentum maps. Finally, Section 7 will present one of the main theorems about symplectic reduction. It concerns the case where the natural configuration space for a system is itself a Lie group G ; (cf. (3) of Section 2.2). Quotienting the natural phase space (the cotangent bundle on G) will give a Poisson manifold that “is” the dual of G ’s Lie algebra.

In this Section, I first sketch some notions of differential geometry, and fix notation (Section 3.1). Then I introduce Lie algebras and Lie brackets of vector fields (Section 3.2). Though most of this Section (indeed this Chapter!) is about differential rather than integral notions, I will later need Frobenius’ theorem, which I present in Section 3.3. Then I give some basic information about Lie groups and their Lie algebras (Section 3.4).

¹⁶And here one should resist being prejudiced because of familiarity. Why *not* have Newtonian gravity arise from a microscopic cyclic degree of freedom? Why *not* have the Lorentz force law arise from geodesic motion in a five-dimensional spacetime with the fifth dimension wrapped up, so that conservation of charge is explained, in Noether’s theorem fashion, by a symmetry?

3.1 Vector fields on manifolds

3.1.1 Manifolds, vectors, curves and derivatives

By way of fixing ideas and notation, I begin by giving details about some ideas in differential geometry (some already used in Section 2.1), and introducing some new notation for them.

A manifold M will be finite-dimensional, except for obvious and explicit exceptions such as the infinite-dimensional group of diffeomorphisms of a (as usual: finite-dimensional!) manifold. I will not be concerned about the degree of differentiability in the definition of a manifold, or of any associated geometric objects: ‘smooth’ can be taken throughout what follows to mean C^∞ . I will often not be concerned with global, as against local, structures and results; (though the reduction results we are driving towards *are* global in nature). For example, I will not be concerned about whether curves are inextendible, or flows are complete.

I shall in general write a vector at a point $x \in M$ as X ; or in terms of local coordinates x^i , as $X = X^i \frac{\partial}{\partial x^i}$ (summation convention). *From now on*, I shall write the tangent space at a point $x \in M$ as $T_x M$ (rather than just T_x), thus explicitly indicating the manifold M to which it is tangent. As before, I write the tangent bundle, consisting of the “meshing collection” of these tangent spaces, as TM . Similarly, I write a 1-form (covector) at a point $x \in M$ as α ; and so the cotangent space at $x \in M$ as $T_x^* M$; and as before, the cotangent bundle as $T^* M$.

A smooth map $f : M \rightarrow N$ between manifolds M and N (maybe $N = M$) maps smooth curves to smooth curves, and so tangent vectors to tangent vectors; and so on for 1-forms and higher tensors. It is convenient to write Tf , called the *derivative* or *tangent* of f (also written as f_* or df or Df), for the induced map on the tangent bundle.

In more detail: let us take a curve c in M to be a smooth map from an interval $I \subset \mathbb{R}$ to M , and a tangent vector at $x \in M$, $X \in T_x M$, to be an equivalence class $[c]_x$ of curves through x . (The equivalence relation is that the curves be tangent at x , with respect to every local chart at x ; but I omit the details of this.) Then we define $Tf : TM \rightarrow TN$ (also written $f_* : TM \rightarrow TN$) by

$$(29) \quad f_*([c]_x) \equiv Tf([c]_x) := [f \circ c]_{f(x)}, \quad \text{for all } x \in M.$$

We sometimes write $T_x f$ for the restriction of Tf to just the tangent space $T_x M$ at x ; i.e.

$$(30) \quad T_x f : [c]_x \in T_x M \mapsto [f \circ c]_{f(x)} \in T_{f(x)} N.$$

In Section 3.1.2.2, we will discuss how one can instead define tangent vectors to be differential operators on the set of all scalar functions defined in some neighbourhood of the point in question, rather than equivalence classes of curves. One can then define the tangent map $f_* \equiv Tf$ in a way provably equivalent to that above.

3.1.2 Vector fields, integral curves and flows

We will be especially concerned with vector fields defined on M , i.e. $X : x \in M \mapsto X(x) \in T_x M$, or on a subset $U \subset M$. So suppose that X is a vector field on M and $f : M \rightarrow N$ is a smooth map, so that $T_x f : T_x M \rightarrow T_{f(x)} N$.

3.1.2.1 Push-forwards and pullbacks It is important to note $(T_x f)(X(x))$ does *not* in general define a vector field on N . For $f(M)$ may not be all of N , so that for $y \in (N - \text{ran}(f))$ $(T_x f)(X(x))$ assigns no element of $T_y N$. And f may not be injective, so that we could have $x, x' \in M$ and $f(x) = f(x')$ with $(T_x f)(X(x)) \neq (T_{x'} f)(X(x'))$. Thus we say that vector fields do not *push forward*.

On the other hand, suppose that $f : M \rightarrow N$ is a *diffeomorphism* onto N : that is, the smooth map f is a bijection, and its inverse f^{-1} is also smooth. Then for any vector field X on M , $Tf(X)$ is a vector field on N . So in this case, the vector field does push forward. Accordingly, $Tf(X)$ is called the *push-forward* of X ; it is often written as $f_*(X)$. So for any $x \in M$, the pushed forward vector field at the image point $f(x)$ is given by

$$(31) \quad (f_*(X))(f(x)) := T_x f \cdot X(x) .$$

(Note the previous use of the asterisk-subscript for the derivative of f , in eq. 29.)

This prompts three more general comments.

1. More generally: we say that two vector fields, X on M and Y on N , are *f-related* on M (respectively: on $S \subset M$) if $(Tf)(X) = Y$ at all $x \in M$ (respectively: $x \in S$).
2. We can generalize the idea that a diffeomorphism implies that a vector field can be pushed forward, in two ways. First, the diffeomorphism need only be defined locally, on some neighbourhood of the point $x \in M$ of interest. Second, a diffeomorphism establishes a one-one correspondence, not just between vector fields defined on its domain and codomain, but also between all differential geometric objects defined on its domain and codomain: in particular, 1-form fields, and higher rank tensors.
3. (This continues comment (2).) Though vector fields do not in general push forward, 1-form fields *do* in general *pull back*. This is written with an asterisk-superscript. That is: for any smooth $f : M \rightarrow N$, not necessarily a diffeomorphism (even locally), and any 1-form field (differential 1-form) α on N , we define the pullback $f^*(\alpha)$ to be the 1-form on M whose action, for each $x \in M$, and each $X \in T_x M$, is given by:

$$(32) \quad (f^*(\alpha))(X) := \alpha|_{f(x)} (Tf(X)) .$$

Similarly, of course if the map f is defined only locally on a subset of M : a 1-form defined on the range of f pulls back to a 1-form on the domain of f .

3.1.2.2 The correspondence between vector fields and flows The leading idea about vector fields is that, for any manifold, the theorems on the local existence, uniqueness and differentiability of solutions of systems of ordinary differential equations (e.g. [Arnold, 1973, 48–49, 77–78, 249–250; Olver, 2000, Prop 1.29]) secure a one-one correspondence between four notions:

- (i) Vector fields X on a subset $U \subset M$, on which they are non-zero; $X : x \in U \mapsto X(x) \in T_x M, X(x) \neq 0$;
- (ii) Non-zero directional derivatives at each point $x \in U$, in the direction of the vector $X(x)$. In terms of coordinates $\mathbf{x} = x^1, \dots, x^n$, these are first-order linear differential operators $X^1(\mathbf{x}) \frac{\partial}{\partial x^1} + \dots + X^n(\mathbf{x}) \frac{\partial}{\partial x^n}$, with $X^i(\mathbf{x})$ the i -component in this coordinate system of the vector $X(x)$. Such an operator is often introduced abstractly as a *derivation*: a map on the set of smooth real-valued functions defined on a neighbourhood of x , that is linear and obeys the Leibniz rule.
- (iii) Integral curves (aka: solution curves) of the fields X in U ; i.e. smooth maps $\phi : I \rightarrow M$ from a real open interval $I \subset \mathbb{R}$ to U , with $0 \in I, \phi(0) = x \in U$, and whose tangent vector at each $\phi(\tau), \tau \in I$ is $X(\phi(\tau))$.
- (iv) Flows X^τ mapping, for each field X and each $x \in U$, some appropriate subset of U to another: $X^\tau : U \rightarrow M$. This flow is guaranteed to exist only in some neighbourhood of a given point x , and for τ in some neighbourhood of $0 \in \mathbb{R}$; but this will be enough for us. Such a flow is a one-parameter subgroup of the “infinite-dimensional group” of all local diffeomorphisms.

I spell out this correspondence in a bit more detail: In local coordinates x^1, \dots, x^n , any smooth curve $\phi : I \rightarrow M$ is given by n smooth functions $\phi(\tau) = (\phi^1(\tau), \dots, \phi^n(\tau))$, and the tangent vector to ϕ at $\phi(\tau) \in M$ is

$$(33) \quad \dot{\phi}(\tau) = \dot{\phi}^1(\tau) \frac{\partial}{\partial x^1} + \dots + \dot{\phi}^n(\tau) \frac{\partial}{\partial x^n}.$$

So for ϕ to be an integral curve of X requires that for all $i = 1, \dots, n$ and all $\tau \in I$

$$(34) \quad \dot{\phi}^i(\tau) = X^i(\tau).$$

The local existence and uniqueness, for a given vector field X and $x \in M$, of the integral curve $\phi_{X,x}$ through x (with $\phi(0) = x$) then ensures that the flow, written either as X^τ or as $\phi_X(\tau)$

$$(35) \quad X^\tau : x \in M \mapsto X^\tau(x) \equiv \phi_{X,x}(\tau) \in M \quad ,$$

is (at least locally) well-defined. The flow is a one-parameter group of transformations of M , and X is called its *infinitesimal generator*.

The exponential notation

$$(36) \quad \exp(\tau X)(x) := X^\tau(x) \equiv \phi_{X,x}(\tau)$$

is suggestive. For example, the group operation in the flow, i.e.

$$(37) \quad X^{\tau+\sigma}(x) = X^\tau(X^\sigma(x)) \quad ,$$

is written in the suggestive notation

$$(38) \quad \exp((\tau + \sigma)X)(x) = \exp(\tau X)(\exp(\sigma X)(x)) \quad .$$

So computing the flow for a given X (i.e. solving a system of n first-order differential equations!) is called *exponentiation* of the vector field X .

Remark: The above correspondence can be related to our discussion of diffeomorphisms and pushing forward vector fields. In particular: if two vector fields, X on M and Y on N , are *f-related* by $f : M \rightarrow N$, so that $(Tf)(X(x)) = Y(f(x))$, then f induces a map from integral curves of X to integral curves of Y . We can express this in terms of exponentiation of X and $Y = (Tf)(X)$:

$$(39) \quad f(\exp(\tau X)x) = \exp(\tau(Tf)(X))(f(x)).$$

Remark: I emphasise that the above correspondence between (i), (ii), (iii) and (iv) is not true at a *single* point. More precisely:

- (a) On the one hand: the correspondence between (i) and (ii) holds at a point; and also holds for zero vectors. That is: a single vector $X \in T_x M$ corresponds to a directional derivative operator (derivation) at x ; and $X = 0$ corresponds to the zero derivative operator mapping all local scalars to 0. (Indeed, as I mentioned: vectors are often defined as such operators/derivations). But:
- (b) On the other hand: the correspondence between (i) and (iii), or between (i) and (iv), requires a neighbourhood. For a single vector $X \in T_x M$ corresponds to a whole class of curves (and so: of flows) through x , not to a single curve. Namely, it corresponds to all the curves (flows) with X as their tangent vector.

However, we shall see (starting in Section 3.4) that for a manifold with suitable extra structure, a single vector *does* determine a curve. (And we will again talk of exponentiation.)

We need to generalize one aspect of the above correspondence (i)–(iv), namely the (i)–(ii) correspondence between vectors and directional derivatives. This generalization is the Lie derivative.

3.1.3 The Lie derivative

Some previous Sections have briefly used the Lie derivative. Since we will use it a lot in the sequel, we now introduce it more thoroughly.

We have seen that given a vector field X on a manifold M , a point $x \in M$, and any scalar function f defined on a neighbourhood of x , there is a naturally defined

rate of change of f along X at x : the directional derivative $X(x)(f)$.

Now we will define the Lie derivative along X as an operator \mathcal{L}_X that defines a rate of change along X : not only for locally defined functions (for which the definition will agree with our previous notion, i.e. we will have $\mathcal{L}_X(f) = X(f)$); but also for vector fields and differential 1-forms.¹⁷ We proceed in three stages.

- (1) We first define the Lie derivative as an operator on scalar functions, in terms of the vector field X on M . We define the *Lie derivative* along the field X (aka: the derivative in the direction of X), \mathcal{L}_X , as the operator on scalar functions $f : M \rightarrow \mathbb{R}$ defined by:

$$(40) \quad \mathcal{L}_X : f \mapsto \mathcal{L}_X f : M \rightarrow \mathbb{R} \text{ with } \forall x \in M : \quad (\mathcal{L}_X f)(x) := \left. \frac{d}{d\tau} \right|_{\tau=0} f(X^\tau(x)) \equiv X(x)(f).$$

Though this definition assumes that both X and f are defined globally, i.e. on all of M , it can of course be restricted to a neighbourhood. Thus defined, \mathcal{L}_X is linear and obeys the Leibniz rule, i.e.

$$(41) \quad \mathcal{L}_X(fg) = f\mathcal{L}_X(g) + g\mathcal{L}_X(f) ;$$

In coordinates $\mathbf{x} = x^1, \dots, x^n$, $\mathcal{L}_X f$ is given by

$$(42) \quad \mathcal{L}_X f = X^1(\mathbf{x}) \frac{\partial f}{\partial x^1} + \dots + X^n(\mathbf{x}) \frac{\partial f}{\partial x^n},$$

with $X^i(\mathbf{x})$ the i -component of the vector $X(x)$. Eq. 42 means that despite eq. 40's mention of the flow X^τ , the Lie derivative of a scalar agrees with our previous notion of directional derivative: that is, for all f , $\mathcal{L}_X(f) = X(f)$.

- (2) In (1), the vector field X determined the operator \mathcal{L}_X : in terms of Section 3.1.2.2's correspondence, we moved from (i) to (ii). But we can conversely define a vector field in terms of its Lie derivative; and in Section 3.2.2's discussion of the Lie bracket, we shall do exactly this.

In a bit more detail: — We note that the set $\mathcal{F}(M)$ of all scalar fields on M , $f : M \rightarrow \mathbb{R}$ forms an (infinite-dimensional) real vector space under pointwise addition. So also does the set $\mathcal{X}(M)$ of all vector fields on M , $X : x \in M \mapsto X(x) \in T_x M$. Furthermore, $\mathcal{X}(M)$ is isomorphic as a real vector space, and as an module over the scalar fields, to the collection of operators \mathcal{L}_X . The isomorphism is given by the map $\theta : X \mapsto \mathcal{L}_X$ defined in (1).

- (3) We now extend the definition of \mathcal{L}_X so as to define it on vector fields Y and 1-forms α . We can temporarily use θ as notation for *either* a vector field Y or a differential 1-form α . Given a vector field X and flow $X^\tau \equiv \phi_X(\tau)$,

¹⁷Indeed, the definition can be extended to all higher rank tensors. But I will not develop those details, since —apart from Section 2.1.3's mention of the Lie derivative of the symplectic form $\mathcal{L}_X \omega$ (viz. the requirement that if X is a symmetry, $\mathcal{L}_X \omega = 0$) —we shall not need them.

we need to compare θ at the point $x \in M$ with θ at the nearby point $X^\tau(x) \equiv \phi_{X,x}(\tau)$, in the limit as τ tends to zero. But the value of θ at $X^\tau(x)$ is in the tangent space, or cotangent space, at $X^\tau(x)$: $T_{X^\tau(x)}M$ or $T_{X^\tau(x)}^*M$. So to make the comparison, we need to somehow transport back this value to T_xM or T_x^*M .

Fortunately, the vector field X provides a natural way to define such a transport. For the vector field Y , we use the differential (i.e. push-forward) of the inverse flow, to “get back” from $X^\tau(x)$ to x . Using $\phi^*(\tau)$ for this “pullback” of $\phi_{X,x}(\tau)$, we define

$$(43) \quad \phi^*(\tau) := T(\exp(-\tau X)) \equiv d \exp(-\tau X) : T_{X^\tau(x)}M \equiv T_{\exp(\tau X)(x)}M \rightarrow T_xM \ .$$

For the 1-form α , we define the transport by the pullback, already defined by eq. 32:

$$(44) \quad \phi^*(\tau) := (\exp(-\tau X))^* : T_{X^\tau(x)}^*M \equiv T_{\exp(\tau X)(x)}^*M \rightarrow T_x^*M \ .$$

With these definitions of $\phi^*(\tau)$, we now define the *Lie derivative* $\mathcal{L}_X\theta$, where θ is a vector field Y or a differential 1-form α , as the vector field or differential 1-form respectively, with value at x given by

$$(45) \quad \lim_{\tau \rightarrow 0} \frac{\phi^*(\tau)(\theta|_{X^\tau(x)}) - \theta|_x}{\tau} = \frac{d}{d\tau} \Big|_{\tau=0} \phi^*(\tau)(\theta|_{X^\tau(x)}) \ .$$

Finally, an incidental result to illustrate this Chapter’s “story so far”. It connects Noether’s theorem, from Section 2.1.3, to this Section’s details about the Lie derivative, and to the theorem stating the local existence and uniqueness of solutions of ordinary differential equations (cf. the start of Section 3.1.2.2). This latter theorem implies that on any manifold any vector field X can be “straightened out”, in the sense that around any point at which X is non-zero, there is a local coordinate system in which X has all but one component vanish and the last component equal to 1. Using this theorem, it is straightforward to show that on any even-dimensional manifold *any* vector field X is locally Hamiltonian, with respect to *some* symplectic form, around a point where X is non-zero. One just defines the symplectic form by Lie-dragging from a surface transverse to X ’s integral curves.

3.2 Lie algebras and brackets

I now introduce Lie algebras and the Lie bracket of two vector fields.

3.2.1 Lie algebras

A *Lie algebra* is a vector space V equipped with a bilinear anti-symmetric operation, usually denoted by square brackets (and called ‘bracket’ or ‘commutator’), $[\cdot, \cdot] : V \times V \rightarrow V$, that satisfies the Jacobi identity, i.e.

$$(46) \quad [[X, Y], Z] + [[Y, Z], X] + [[Z, X], Y] = 0 \ .$$

3.2.1.1 Examples; rotations introduced

Here are three examples.

- (i) $n \times n$ matrices equipped with the usual commutator, i.e. $[X, Y] := XY - YX$. (So the matrix multiplication “contributes” to the bracket, but not to the underlying vector space structure.)
- (ii) 3×3 anti-symmetric matrices, equipped with the usual commutator.
- (iii) \mathbb{R}^3 equipped with vector multiplication.

In fact, example (iii) is essentially the same as example (ii); and this example will recur in what follows, in connection with rotations and the rigid body. (We will also see that example (ii) is in a sense more fundamental.)

To explain this, we first recall that every anti-symmetric operator A on a three-dimensional oriented euclidean space is the operator of vector multiplication by a fixed vector, ω say. That is: for all \mathbf{q} , $A\mathbf{q} = [\omega, \mathbf{q}] \equiv \omega \wedge \mathbf{q}$. (Proof: the anti-symmetric operators on \mathbb{R}^3 for a 3-dimensional vector space, since an anti-symmetric 3×3 matrix has three independent components. Vector multiplication by a vector ω is a linear and anti-symmetric operator; varying ω we get a subspace of the space of all anti-symmetric operators on \mathbb{R}^3 ; but this subspace has dimension 3; so it coincides with the space of all anti-symmetric operators.)

With this result in hand, the following three points are all readily verified.

- (1) The matrix representation of A in cartesian coordinates is then

$$(47) \quad A = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix}.$$

We can write

$$(48) \quad A \leftrightarrow \omega \quad \text{or} \quad A_{ij} = -\epsilon_{ijk}\omega_k \quad \text{or} \quad \omega_i = -\frac{1}{2}\epsilon_{ijk}A_{jk}.$$

- (2) The plane Π of vectors perpendicular to ω is an invariant subspace for A , i.e. $A(\Pi) = \Pi$. And ω is an eigenvector for A with eigenvalue 0. This suggests a familiar elementary interpretation, which will be confirmed later (Section 3.4): viz. that any 3×3 anti-symmetric matrix A represents a infinitesimal rotation, and ω represents instantaneous angular velocity. That is, we will have, for all $\mathbf{q} \in \mathbb{R}^3$: $\dot{\mathbf{q}} = A\mathbf{q} = [\omega, \mathbf{q}]$.
- (3) The commutator of any two 3×3 anti-symmetric matrices A, B , i.e. $[A, B] := AB - BA$, corresponds by eq. 48 to vector multiplication of the axes of rotation. That is: writing eq. 48's bijection from vectors to matrices as $\Theta : \omega \mapsto A =: \Theta(\omega)$, we have for vectors $\mathbf{q}, \mathbf{r}, \mathbf{s}$

$$(49) \quad (\Theta(\mathbf{q})\Theta(\mathbf{r}) - \Theta(\mathbf{r})\Theta(\mathbf{q}))\mathbf{s} = \Theta(\mathbf{q})[\mathbf{r}, \mathbf{s}] - \Theta(\mathbf{r})[\mathbf{q}, \mathbf{s}]$$

$$(50) \quad \quad \quad = [\mathbf{q}, [\mathbf{r}, \mathbf{s}]] - [\mathbf{r}, [\mathbf{q}, \mathbf{s}]]$$

$$(51) \quad \quad \quad = [[\mathbf{q}, \mathbf{r}], \mathbf{s}] = \Theta([\mathbf{q}, \mathbf{r}]) \cdot \mathbf{s}.$$

where the $[\cdot]$ represents vector multiplication, i.e. $[\mathbf{q}, \mathbf{r}] \equiv \mathbf{q} \wedge \mathbf{r}$.

Eq. 51 means that Θ gives a Lie algebra isomorphism; and so our example (iii) is essentially the same as example (ii).

Besides, we can already glimpse why example (ii) is in a sense more fundamental. For this correspondence between anti-symmetric operators (or matrices) and vectors, eq. 48, is specific to three dimensions. In n dimensions, the number of independent components of an anti-symmetric matrix is $n(n-1)/2$: only for $n = 3$ is this equal to n . Yet we will see later (Section 3.4.4) that rotations on euclidean space \mathbb{R}^n of any dimension n are generated, in a precise sense, by the Lie algebra of $n \times n$ anti-symmetric matrices. So only for $n = 3$ is there a corresponding representation of rotations by vectors in \mathbb{R}^n .

In the next two Subsections, we shall see other examples of Lie algebras: whose vectors are vector fields (Section 3.2.2), or tangent vectors at the identity element of a Lie group (Section 3.4). The first example will be an infinite-dimensional Lie algebra; the second finite-dimensional (since we will only consider finite-dimensional Lie groups). Besides, the above examples (i) and (ii) (equivalently: (i) and (iii)) will recur: each will be the vector space of tangent vectors at the identity element of a Lie group.

3.2.1.2 Structure constants A finite-dimensional Lie algebra is characterized, relative to a basis, by a set of numbers, called *structure constants* that specify the bracket operation. Thus if $\{v_1, \dots, v_n\}$ is a basis of a Lie algebra V , we define the structure constants c_{ij}^k , ($i, j, k = 1, \dots, n$) by expanding, in terms of this basis, the bracket of any two basis elements

$$(52) \quad [v_i, v_j] = \sum_k c_{ij}^k v_k .$$

The bilinearity of the bracket implies that eq. 52 determines the bracket of all pairs of vectors $v, w \in V$. And the bracket's obeying anti-symmetry and the Jacobi identity implies that, for any basis, the structure constants obey

$$(53) \quad c_{ij}^k = -c_{ji}^k ; \quad \sum_k (c_{ij}^k c_{kl}^m + c_{li}^k c_{kj}^m + c_{jl}^k c_{ki}^m) = 0.$$

Conversely, any set of constants c_{ij}^k obeying eq. 53 are the structure constants of an n -dimensional Lie algebra.

3.2.2 The Lie bracket of two vector fields

Given two vector fields X, Y on a manifold M , the corresponding flows do not in general commute: $X^t Y^s \neq Y^s X^t$. The non-commutativity is measured by the commutator of the Lie derivatives of X and of Y , i.e. $\mathcal{L}_X \mathcal{L}_Y - \mathcal{L}_Y \mathcal{L}_X$. (Cf. eq. 40 and 45 for a definition of the Lie derivative.) Here, 'measured' can be made precise by considering Taylor expansions; but I shall not go into detail about this.

What matters for us is that this commutator, which is at first glance seems to be a second-order operator, is in fact a first-order operator. This is verified by

calculating in a coordinate system, and seeing that the second derivatives occur twice with opposite signs:

$$(54) (\mathcal{L}_X \mathcal{L}_Y - \mathcal{L}_Y \mathcal{L}_X) f = \Sigma_i X^i \frac{\partial}{\partial x^i} \left(\Sigma_j Y^j \frac{\partial f}{\partial x^j} \right) - \Sigma_j Y^j \frac{\partial}{\partial x^j} \left(\Sigma_i X^i \frac{\partial f}{\partial x^i} \right)$$

$$(55) \quad = \dots = \Sigma_{i,j} \left(X^i \frac{\partial Y^j}{\partial x^i} - Y^j \frac{\partial X^i}{\partial x^j} \right) \frac{\partial f}{\partial x^j}.$$

So $\mathcal{L}_X \mathcal{L}_Y - \mathcal{L}_Y \mathcal{L}_X$ corresponds to a vector field: (recall (2) of Section 3.1.3, about defining a vector field from its Lie derivative). We call this field Z the *Lie bracket* (also known as: Poisson bracket, commutator, and Jacobi-Lie bracket!) of the fields X and Y , and write it as $[X, Y]$. It is also written as $\mathcal{L}_X Y$ and called the *Lie derivative of Y with respect to X* . (Beware: some books use an opposite sign convention.)

Thus $Z \equiv [X, Y] \equiv \mathcal{L}_X Y$ is defined to be the vector field such that

$$(56) \quad \mathcal{L}_Z \equiv \mathcal{L}_{[X, Y]} = \mathcal{L}_X \mathcal{L}_Y - \mathcal{L}_Y \mathcal{L}_X .$$

It follows that $Z \equiv [X, Y]$'s components in a coordinate system are given by eq. 55. This formula can be remembered by writing it (with summation convention, i.e. omitting the Σ) as

$$(57) \quad \left[X^i \frac{\partial}{\partial x^i}, Y^j \frac{\partial}{\partial x^j} \right] = X^i \frac{\partial Y^j}{\partial x^i} \frac{\partial}{\partial x^j} - Y^j \frac{\partial X^i}{\partial x^j} \frac{\partial}{\partial x^i}$$

Another way to write eq. 55 is as:

$$(58) \quad [X, Y]^j = (X \cdot \nabla) Y^j - (Y \cdot \nabla) X^j ;$$

or without coordinates, writing \mathbf{D} for the derivative map given by the Jacobian matrix, as

$$(59) \quad [X, Y] = \mathbf{D}Y \cdot X - \mathbf{D}X \cdot Y.$$

Again, the vector field $Z \equiv [X, Y]$ measures the non-commutation of the flows X^t and Y^s : in particular, these flows commute iff $[X, Y] = 0$.

We will need three results about the Lie bracket. They concern, respectively, the relation to Lie algebras, to Poisson brackets, and to Frobenius' theorem.

- (1) The Lie bracket is obviously a bilinear and anti-symmetric operation on the (infinite-dimensional) vector space $\mathcal{X}(M)$ of all vector fields on M : $[\cdot, \cdot] : \mathcal{X}(M) \times \mathcal{X}(M) \rightarrow \mathcal{X}(M)$. One readily checks that it satisfied the Jacobi identity. (Expand $\mathcal{L}_{[[X, Y], Z]} = \mathcal{L}_{[X, Y]} \mathcal{L}_Z - \mathcal{L}_Z \mathcal{L}_{[X, Y]}$ etc.) So: $\mathcal{X}(M)$ is an (infinite-dimensional) Lie algebra.
- (2) Returning to Hamiltonian mechanics (Section 2.1): there is a simple and fundamental relation between the Lie bracket and the Poisson bracket, via the notion of Hamiltonian vector fields (Section 2.1.3). Namely: the Hamiltonian vector field of the Poisson bracket of two scalar functions f, g on the symplectic manifold M is, upto a sign, the Lie bracket of the Hamiltonian vector fields, X_f and X_g , of f and g :

$$(60) \quad X_{\{f,g\}} = -[X_f, X_g] = [X_g, X_f].$$

Proof: apply the rhs to an arbitrary scalar $h : M \rightarrow \mathbb{R}$. One easily obtains $X_{\{f,g\}}(h)$, by using:

- (i) the definition of a Hamiltonian vector field;
- (ii) the Lie derivative of a function equals its elementary directional derivative eq. 40; and
- (iii) the Poisson bracket is antisymmetric and obeys the Jacobi identity.

This result means that the Hamiltonian vector fields on a symplectic manifold M , equipped with the Poisson bracket, form an (infinite-dimensional) Lie subalgebra of the Lie algebra $\mathcal{X}(M)$ of all vector fields on the symplectic manifold M . Later, it will be important that this result extends from symplectic manifolds to Poisson manifolds; (details in Section 5.2.2).

- (3) For Frobenius' theorem (Section 3.3), we need to relate the Lie bracket to Section 3.1.2's idea of vector fields being f -related by a map $f : M \rightarrow N$ between manifolds M and N . In short: if two pairs of vector fields are f -related, so is their Lie bracket. More explicitly: if X, Y are vector fields on M , and $f : M \rightarrow N$ is a map such that $(Tf)(X), (Tf)(Y)$ are well-defined vector fields on N , then Tf commutes with the Lie bracket:

$$(61) \quad (Tf)[X, Y] = [(Tf)X, (Tf)Y].$$

3.3 Submanifolds and Frobenius' theorem

This Subsection differs from the preceding ones in three ways. First, it emphasises integral, rather than differential, notions.

Second: Section 3.1.2.2 have emphasised that the integral curves of a vector field correspond to integrating a system of ordinary differential equations. Since such curves are one-dimensional submanifolds of the given manifold, our present topic, viz. higher-dimensional submanifolds, naturally suggests partial differential equations. For their integration involves finding, given an assignment to each point x of a manifold M of a subspace S_x (with dimension greater than one) of the tangent space $T_x M$, an integral surface, i.e. a submanifold S of M whose tangent space at each of its points is S_x .¹⁸

However, we will *not* be concerned with partial differential equations. For us, submanifolds of dimension higher than one arise when the span S_x of the tangent vectors at x to a set of *vector* fields fit together to form a submanifold. Thus Frobenius' theorem states, roughly speaking, that a finite set of vector fields is integrable in this sense iff the vector fields are in involution. That is: iff their

¹⁸Beware: there is no analogue for partial differential equations of the local existence and uniqueness theorem for ordinary differential equations. Even a field of two-dimensional planes in three-dimensional space is in general not integrable, e.g. the field of planes given by the equation $dz = ydx$. So integrable fields of planes, or other tangent subspaces on a manifold, are an exception; and accordingly, the integration theory for partial differential equations is less unified, and more complicated, than that for ordinary differential equations.

pairwise Lie brackets are expandable in terms of the fields; i.e. the vector fields form a Lie subalgebra of the entire Lie algebra of vector fields. We will not need to prove this theorem. But we need to state it and use it—in particular, for the foliation of Poisson manifolds.

Third: a warning is in order. The intuitive idea of a subset $S \subset M$ that is a smooth manifold “in its own right” can be made precise in different ways. So there are subtleties about the definition of ‘submanifold’, and terminology varies between expositions—in a way it does not for the material in previous Sections. I will adopt what seems to be a widespread, if not majority, terminology.¹⁹

3.3.1 Submanifolds

The fundamental definition is:

Given a manifold M ($\dim(M)=n$), a *submanifold* of M of dimension k is a subset $N \subset M$ such that for every $y \in N$ there is an admissible local chart (i.e. a chart in M 's maximal atlas) (U, ϕ) with $y \in U$ and with the *submanifold property*, viz.

$$(62) \text{ (SM). } \phi : U \rightarrow \mathbb{R}^k \times \mathbb{R}^{n-k} \text{ and } \phi(U \cap N) = \phi(U) \cap (\mathbb{R}^k \times \{\mathbf{0}\}).$$

The set N becomes a manifold, generated by the atlas of all charts of the form $(U \cap N, \phi | (U \cap N))$, where (U, ϕ) is a chart of M having the submanifold property. (This makes the topology of N the relative topology.)

We need to take note of two ways in which submanifolds can be specified in terms of smooth functions between manifolds.

- (1) A submanifold can be specified as the set on which a smooth function $f : M \rightarrow P$ between manifolds takes a certain value. In effect, this will be a generalization of eq. 62's requirement that $n - k$ coordinate-components of a chart ϕ take the value zero. This will involve the idea that the tangent map Tf is surjective, in which case f will be called a *submersion*. We will need this approach for quotients of actions of Lie groups.
- (2) A submanifold can be specified parametrically, as the set of values of a local parametrization: i.e. as the range of a smooth function f with M as *codomain*. This will involve the idea that the tangent map Tf is injective, in which case f will be called an *immersion*. We will need this approach for Frobenius' theorem.

¹⁹My treatment is based on [Marsden and Ratiu, 1999, p. 124–127, 140] for Section 3.3.1, and [Olver, 2000, p. 38–40] for Section 3.3.2. As to varying terminology: Olver [2000, p. 9] defines ‘submanifold’ to be what we will call an immersed submanifold; (which latter, for us, does *not* have to be a submanifold, since the immersion need not be an embedding). Bishop and Goldberg [1980, p. 40–41] provide a similar example. For a detailed introduction to the different notions of submanifold, cf. Darling [1994, Chapters 3 and 5]. Note that I will also omit some details, in particular about Frobenius' theorem providing regular immersions.

(1): *Submersions*:

If $f : M \rightarrow P$ is a smooth map between manifolds, a point $x \in M$ is called a *regular point* if the tangent map $T_x f$ is surjective; otherwise x is a *critical point* of f . If $C \subset M$ is the set of critical points of M , we say $f(C)$ is the set of *critical values* of f , and $P - f(C)$ is the set of *regular values* of f . So if $p \in P$ is a regular value of f , then at every $x \in M$ with $f(x) = p$, $T_x f$ is surjective.

The *submersion theorem* states that if $p \in P$ is a regular value of f , then:

- (i) $f^{-1}(p)$ is a submanifold of M of dimension $\dim(M) - \dim(P)$; and
- (ii) the tangent space of this submanifold at any point $x \in f^{-1}(p)$ is the kernel of f 's tangent map:

$$(63) \quad T_x(f^{-1}(p)) = \ker T_x f .$$

If $T_x f$ is surjective for every $x \in M$, f is called a *submersion*.

(2): *Immersion*s:

A smooth map between manifolds $f : M \rightarrow P$ is called an *immersion* if $T_x f$ is injective at every $x \in M$. The *immersion theorem* states that $T_x f$ is injective iff there is a neighbourhood U of x in M such that $f(U)$ is a submanifold of P and $f|_U : U \rightarrow f(U)$ is a diffeomorphism.

NB: This does not say that $f(M)$ is a submanifold of P . For f may not be injective (so that $f(M)$ has self-intersections). And even if f is injective, f can fail to be a homeomorphism between M and $f(M)$, equipped with the relative topology induced from P . A standard simple example is an injection of an open interval of \mathbb{R} into an ‘almost-closed’ figure-of-eight in \mathbb{R}^2 .

Nevertheless, when $f : M \rightarrow P$ is an immersion, and is also injective, we call $f(M)$ an *injectively immersed submanifold* (or shorter: an *immersed submanifold*): though $f(M)$ might not be a submanifold.

We also define an *embedding* to be an immersion that is also a homeomorphism (and so injective) between M and $f(M)$ (where the latter has the relative topology induced from P). If f is an embedding, $f(M)$ is a submanifold of N and f is a diffeomorphism $f : M \rightarrow f(M)$.

In fact, Frobenius’ theorem will provide injectively immersed submanifolds that need not be embedded, and so need not be submanifolds. (They must also obey another condition, called ‘regularity’, that I will not go into.)

3.3.2 The theorem

We saw at the end of Section 3.2.2 that if two pairs of vector fields are f -related, so is their Lie bracket: cf. eq. 61. This result immediately yields a necessary condition for two vector fields to be tangent to an embedded submanifold: namely

If X_1, X_2 are vector fields on M that are tangent to an embedded submanifold S (i.e. at each $x \in S$, $X_i(x) \in T_x S \subset T_x M$), then their Lie bracket $[X_1, X_2]$ is also tangent to S .

This follows by considering the diffeomorphism $f : \tilde{S} \rightarrow S$ that gives an embedding of S in M . One then uses the fact that Tf commutes with the Lie bracket, eq. 61. That is: the Lie bracket of the f -related vector fields \tilde{X}_1, \tilde{X}_2 on \tilde{S} , which is of course tangent to \tilde{S} , is carried by Tf to the Lie bracket $[X_1, X_2]$ of X_1 and X_2 . So $[X_1, X_2]$ is tangent to S .

The idea of Frobenius' theorem will be that this necessary condition of two vector fields being tangent to a submanifold is also sufficient. To be more precise, we need the following definitions.

A *distribution* D on a manifold M is a subset of the tangent bundle TM such that at each $x \in M$, $D_x := D \cap T_x M$ is a vector space. The dimension of D_x is the *rank* of D at x . If the rank of D is constant on M , we say the distribution is *regular*.

A distribution is *smooth* if for every $x \in M$, and every $X_0 \in D_x$, there is a neighbourhood $U \subset M$ of x , and a smooth vector field X on U such that (i) $X(x) = X_0$, (ii) for all $y \in U$, $X(y) \in D_y$. Such a vector field X is called a *local section* of D . Example: a set of r vector fields, X_1, \dots, X_r each defined on M , together define a smooth distribution of rank at most r .

A distribution is *involutive* if for any pair X_1, X_2 of local sections, the Lie bracket $[X_1, X_2](y) \in D_y$ in the two sections' common domain of definition.

We similarly say that a set of r smooth vector fields, X_1, \dots, X_r , on a manifold M is *in involution* if everywhere in M they span their Lie brackets. That is: there are smooth real functions $h_{ij}^k : M \rightarrow R, i, j, k = 1, \dots, r$ such that at each $x \in M$

$$(64) \quad [X_i, X_j](x) = \sum_k h_{ij}^k(x) X_k(x).$$

(Beware: involution is used in a different sense in connection with Liouville's theorem, viz. a set of real functions on phase space is said to be in involution when all their pairwise Poisson brackets vanish.)

A distribution D on M is *integrable* if for each $x \in M$ there is a local submanifold $N(x)$ of M whose tangent bundle equals the restriction of D to $N(x)$. If D is integrable, the various $N(x)$ can be extended to get, through each $x \in M$, a unique maximal connected set whose tangent space at each of its elements y is D_y . Such a set is called a (maximal) *integral manifold*.

NB: In general, each integral manifold is injectively immersed in M , but not embedded in it; and so, by the discussion in (2) of Section 3.3.1, an integral manifold might not be a submanifold of M . But (like most treatments), I shall ignore this point, and talk of them as submanifolds, *integral submanifolds*.

If the rank of D is constant on M , all the integral submanifolds have a common dimension: the rank of D . But in general the rank of D varies across M , and so does the dimension of the integral submanifolds.

We similarly say that a set of r vector fields, X_1, \dots, X_r , is *integrable*; viz. if through every $x \in M$ there passes a local submanifold $N(x)$ of M whose tangent space at each of its points is spanned by X_1, \dots, X_r . (Again: we allow that at some x , $X_1(x), \dots, X_r(x)$ may be linearly dependent, so that the dimension of the submanifolds varies.)

We say (both for distributions and sets of vector fields) that the collection of integral manifolds is a *foliation* of M , and its elements are *leaves*. Again: if the dimension of the leaves is constant on M , we say the foliation is *regular*.

With these definitions in hand, we can now state Frobenius' theorem: both in its usual form, which concerns the case of constant rank, i.e. regular distributions and vector fields that are everywhere linearly independent; and in a generalized form. The usual form is:

Frobenius' theorem (usual form) A smooth regular distribution is integrable iff it is involutive.

Or in terms of vector fields: a set of r smooth vector fields, X_1, \dots, X_r , on a manifold M , that are everywhere linearly independent, is integrable iff it is in involution.

The generalization comes in two stages. The first stage concerns varying rank, but assumes a finite set of vector fields. It is straightforward: this very same statement holds. That is: a set of r smooth vector fields, X_1, \dots, X_r , on a manifold M (perhaps not everywhere linearly independent) is integrable iff it is in involution.

But for the foliation of Poisson manifolds (Section 5.3.3), we need to consider an infinite set of vector fields, perhaps with varying rank; and for such a set, this statement fails. Fortunately, there is a useful generalization; as follows.

Let \mathcal{X} be a set of vector fields on a manifold M , that forms a vector space. So in the above discussion of r vector fields, \mathcal{X} can be taken as all the linear combinations $\sum_{i=1}^r f_i(x)X_i(x)$, $x \in M$, where the f_i are arbitrary smooth functions $f : M \rightarrow \mathbb{R}$. Such an \mathcal{X} is called *finitely generated*.

For any \mathcal{X} forming a vector space, we say (as before) that \mathcal{X} is *in involution* if $[X, Y] \in \mathcal{X}$ whenever $X, Y \in \mathcal{X}$. Let \mathcal{X}_x be the subspace of $T_x M$ spanned by the $X(x)$ for all $X \in \mathcal{X}$. As before, we define: an *integral manifold* of \mathcal{X} is a submanifold $N \subset M$ such that for all $y \in N$, $T_y N = \mathcal{X}_y$; and \mathcal{X} is called *integrable* iff through each $x \in M$ there passes an integral manifold.

As before: if \mathcal{X} is integrable, it is in involution. But the converse fails. A further condition is needed, as follows.

We say that \mathcal{X} is *rank-invariant* if for any vector field $X \in \mathcal{X}$, the dimension of the subspace $\mathcal{X}_{\exp(\tau X)(x)}$ along the flow generated by X is a constant, independent of τ . (But it can depend on the point x .)

Since the integral curve $\exp(\tau X)(x)$ through x should be contained in any integral submanifold, rank-invariance is certainly a necessary condition of integrability. (It also follows from \mathcal{X} being finitely generated.) In fact we have:

Frobenius' theorem (generalized form) A system \mathcal{X} of vector fields on M is integrable iff it is rank-invariant and in involution.

The idea of the proof is to directly construct the integral submanifolds. The submanifold through x is obtained as

$$(65) \quad N = \{\exp(X_1) \exp(X_2) \dots \exp(X_p)(x) : p \geq 1, X_i \in \mathcal{X}\}.$$

The rank-invariance secures that for any $y \in N$, \mathcal{X}_y has dimension $\dim(N)$.

3.4 Lie groups, and their Lie algebras

I introduce Lie groups and their Lie algebras. By the last two Subsections (Sections 3.4.3 and 3.4.4), we will have enough theory to compute efficiently the Lie algebra of a fundamentally important Lie group, the rotation group.

3.4.1 Lie groups and matrix Lie groups

A *Lie group* is a group G which is also a manifold, and for which the product and inverse operations $G \times G \rightarrow G$ and $G \rightarrow G$ are smooth.

Examples:

- (i) \mathbb{R}^n under addition.
- (ii) The group of linear isomorphisms of \mathbb{R}^n to \mathbb{R}^n , denoted $GL(n, \mathbb{R})$ and called the *general linear group*; represented by the real invertible $n \times n$ matrices. This is an open subset of \mathbb{R}^{n^2} , and so a manifold of dimension n^2 ; and the formulas for the product and inverse of matrices are smooth in the matrix components.
- (iii) The group of rotations about the origin of \mathbb{R}^3 , represented by 3×3 orthogonal matrices of determinant 1; denoted $SO(3)$, where S stands for ‘special’ (i.e. determinant 1), and O for ‘orthogonal’.

In fact, all three examples can be regarded as Lie groups of matrices, with matrix multiplication as the operation. In example (i), consider the isomorphism θ between \mathbb{R}^n under addition and $(n+1) \times (n+1)$ matrices with diagonal entries all equal to 1, other rightmost column entries equal to the given vector in \mathbb{R}^n , and all other entries zero. Thus consider, for the case $n = 3$:

$$(66) \quad \theta : \begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} 1 & 0 & 0 & x \\ 0 & 1 & 0 & y \\ 0 & 0 & 1 & z \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

This suggests that we define a *matrix Lie group* to be any set of invertible real matrices, under matrix multiplication, that is closed under multiplication, inversion and taking of limits. That a matrix Lie group is a Lie group will then follow from $GL(n, \mathbb{R})$ being a Lie group, and the theorem below (in Section 3.4.3) that any closed subgroup of a Lie group is itself a Lie group.

For matrix Lie groups, some of the theory below simplifies. For example, the definition of exponentiation of an element of the group’s Lie algebra reduces to exponentiation of a matrix. But we will develop some of the general theory, since (as always!) it is enlightening and powerful.

3.4.2 The Lie algebra of a Lie group

The main result in this Subsection is that for any Lie group G , the tangent space $T_e G$ at the identity $e \in G$ has a natural Lie algebra structure that is induced by certain natural vector fields on G ; as follows.

3.4.2.1 Left-invariant vector fields define the Lie algebra :

Let G be a Lie group. Each $g \in G$ defines a diffeomorphism of G onto itself by *left translation*, and similarly by *right translation*:

$$(67) \quad L_g : h \in G \mapsto gh \in G \quad ; \quad R_g : h \in G \mapsto hg \in G.$$

Remark: In Section 4 we will describe this in the language of group actions, saying that in eq. 67 G acts on itself by left and right translation.

Now consider the induced maps on the tangent spaces, i.e. the tangent (aka: derivative) maps; cf. eq.s 29, 30. They are $(L_g)_* =: L_{g*}, (R_g)_* =: R_{g*}$ where for each $h \in G$:

$$(68) \quad L_{g*} : T_h G \rightarrow T_{gh} G \quad \text{and} \quad R_{g*} : T_h G \rightarrow T_{hg} G.$$

In particular: the derivative $(R_g)_*$ at $e \in G$ maps $T_e G$ to $T_g G$. This implies that every vector $\xi \in T_e G$ defines a vector field on G : its value at any $g \in G$ is the image $(R_g)_* \xi$ of ξ under $(R_g)_*$. Such a vector field is called a *right-invariant* vector field: it is uniquely defined by (applying the derivative of right translation to) its value at the identity $e \in G$.

In more detail, and now defining *left-invariant* vector fields: —

A vector field X on G is called *left-invariant* if for every $g \in G$, $(L_g)_* X = X$. More explicitly, let us write $T_h L_g$ for the tangent or derivative of L_g at h , i.e. for $L_{g*} : T_h G \rightarrow T_{gh} G$. Then left-invariance requires that

$$(69) \quad (T_h L_g)X(h) = X(gh) \quad \text{for every } g \text{ and } h \in G.$$

Thus every vector $\xi \in T_e G$ defines a left-invariant vector field, written X_ξ , on G : X_ξ 's value at any $g \in G$ is the image $(L_g)_* \xi$ of ξ under $(L_g)_*$. In other words: $X_\xi(g) := (T_e L_g)\xi$.

Not only is a left-invariant vector field uniquely defined by its value at the identity $e \in G$. Also, the set $\mathcal{X}_L(G)$ of left-invariant vector fields on G is isomorphic as a vector space to the tangent space $T_e G$ at the identity e . For the linear maps α, β defined by

$$(70) \quad \alpha : X \in \mathcal{X}_L(G) \mapsto X(e) \in T_e G \quad ; \quad \text{and} \quad \beta : \xi \in T_e G \mapsto \{g \mapsto X_\xi(g) := (T_e L_g)\xi\} \in \mathcal{X}_L(G)$$

compose to give the identity maps:

$$(71) \quad \beta \circ \alpha = id_{\mathcal{X}_L(G)} \quad ; \quad \alpha \circ \beta = id_{T_e G}.$$

$\mathcal{X}_L(G)$ is a Lie subalgebra of the Lie algebra of all vector fields on G , because it is closed under the Lie bracket. That is: the Lie bracket of left-invariant vector fields X and Y is itself left-invariant, since one can check that for every $g \in G$ we have (with L meaning ‘left’ not ‘Lie’!)

$$(72) \quad L_{g*}[X, Y] = [L_{g*}X, L_{g*}Y] = [X, Y].$$

If we now define a bracket on T_eG by

$$(73) \quad [\xi, \eta] := [X_\xi, X_\eta](e)$$

then T_eG becomes a Lie algebra. It is called *the Lie algebra of G* , written \mathfrak{g} (or, to avoid ambiguity about which Lie group is in question: $\mathfrak{g}(G)$). It follows from eq. 72 that

$$(74) \quad [X_\xi, X_\eta] = X_{[\xi, \eta]} \quad ;$$

that is to say, the maps α, β are Lie algebra isomorphisms.

This result, that T_eG has a natural Lie algebra structure, is *very* important. For, as we shall see in the rest of Section 3.4: the structure of a Lie group is very largely determined by the structure of this Lie algebra. Accordingly, as we shall see in Sections 4 and 5 et seq.: this Lie algebra underpins most of the constructions made with the Lie group, e.g. in Lie group actions. Thus Olver writes that this result ‘is the cornerstone of Lie group theory ... almost the entire range of applications of Lie groups to differential equations ultimately rests on this one construction!’ [Olver, 2000, 42].

Before turning in the next Subsection to examples, and the topic of subgroups and subalgebras, I end with four results, (1)-(4), which will be needed later; and a remark.

3.4.2.2 Four results

- (1) Lie group structure determines Lie algebra structure in the following sense. If G, H are Lie groups, and $f : G \rightarrow H$ is a smooth homomorphism, then the derivative of f at the identity $T_e f : \mathfrak{g}(G) \rightarrow \mathfrak{g}(H)$ is a Lie algebra homomorphism. In particular, for all $\xi, \eta \in \mathfrak{g}(G)$, $(T_e f)[\xi, \eta] = [T_e f(\xi), T_e f(\eta)]$. (Cf. eq. 61.)

- (2) *Exponentiation again; a correspondence between left-invariant vector fields and one-dimensional subgroups:*

Recall from Section 3.1, especially eq. 36, that each vector field X on the manifold G determines an integral curve ϕ_X in G passing through the identity e (with $\phi_X(0) = e$). We now write the points in (the image of) this curve as g_τ (X and e being understood):

$$(75) \quad \exp(\tau X)(e) \equiv X^\tau(e) \equiv \phi_{X,e}(\tau) =: g_\tau.$$

It is straightforward to show that if X is left-invariant, this (image of a) curve is a *one-parameter subgroup* of G : i.e. not just as eq. 35 et seq., a one-parameter subgroup of the group of diffeomorphisms of the manifold G . In fact:

$$(76) \quad g_{\tau+\sigma} = g_\tau g_\sigma \quad g_0 = e \quad g_\tau^{-1} = g_{-\tau} .$$

Besides, the group is defined for all $\tau \in \mathbb{R}$; and is isomorphic to either \mathbb{R} or the circle group $SO(2)$. Conversely, any connected one-parameter subgroup of G is generated by a left-invariant vector field in this way.

Accordingly, we define exponentiation of elements ξ of \mathfrak{g} by reference to the isomorphisms eq. 70 and 71. It is also convenient to define this as a map taking values in G . Thus for $\xi \in \mathfrak{g}$ and its corresponding left-invariant vector field X_ξ that takes as value at $g \in G$, $X_\xi(g) := (T_e L_g)(\xi)$, we write the integral curve of X_ξ that passes through e (with value e for argument $\tau = 0$) as

$$(77) \quad \phi_\xi : \tau \in \mathbb{R} \mapsto \exp(\tau X_\xi)(e) \in G .$$

Then we define the *exponential map* of \mathfrak{g} into G to be the map

$$(78) \quad \exp : \xi \in \mathfrak{g} \mapsto \phi_\xi(1) \in G .$$

Using the linearity of β as defined by eq. 70, these two equations, eq. 77 and 78, are related very simply:

$$(79) \quad \exp(\tau \xi) := \phi_{\tau \xi}(1) := \exp(1 \cdot X_{\tau \xi})(e) = \exp(\tau X_\xi) .$$

We write \exp_G rather than \exp when the context could suggest a Lie group other than G .

The map \exp is a local diffeomorphism of a neighbourhood of $0 \in \mathfrak{g}$ to a neighbourhood of $e \in G$; but not in general a global diffeomorphism onto G . In modern terms, this result follows by applying the inverse function theorem to the discussion above. (It also represents an interesting example of the history of subject; cf. [Hawkins, 2000, 82–83] for Lie's version of this result, without explicit mention of its local nature.)

The map \exp also has the basic property, adding to result (1) above, that ...

(3) *Homomorphisms respect exponentiation:*

If $f : G \rightarrow H$ is a smooth homomorphism of Lie groups, then for all $\xi \in \mathfrak{g}$,

$$(80) \quad f(\exp_G \xi) = \exp_H((T_e f)(\xi)) .$$

(4) *Right-invariant vector fields as an alternative approach:*

We have followed the usual practice of defining \mathfrak{g} in terms of left-invariant vector fields. One can instead use right-invariant vector fields. This produces some changes in signs, and in whether certain defined operations respect or

reverse the order of two elements used in their definition. I will not go into many details about this. But some will be needed when we consider:

- (i) Lie group actions, and especially their infinitesimal generators (Section 4.4 and 4.5);
- (ii) reduction on the cotangent bundle of a Lie group —as occurs in the theory of the rigid body (Section 6.5 and 7.3.3).

For the moment we just note two basic results, (A) and (B); postponing others to Section 4.4 et seq.

- (A) Corresponding to the vector space isomorphism between \mathfrak{g} and the left-invariant vector fields, as in eq. 70. viz.

$$(81) \quad \xi \in T_e G \mapsto X_\xi \in \mathcal{X}_L(G) \quad \text{with} \quad X_\xi(g) := (T_e L_g)\xi ,$$

there is a vector space isomorphism to the set of right-invariant vector fields

$$(82) \quad \xi \in T_e G \mapsto Y_\xi \in \mathcal{X}_R(G) \quad \text{with} \quad Y_\xi(g) := (T_e R_g)\xi .$$

Besides, the Lie bracket of right-invariant vector fields is itself right-invariant. So corresponding to our previous definition, eq. 73, of a Lie bracket on $T_e G$, and its corollary eq. 74, i.e. $[X_\xi, X_\eta] = X_{[\xi, \eta]}$, that makes $T_e G \cong \mathcal{X}_L(G)$ a Lie algebra isomorphism: we can also define a Lie bracket on $T_e G$ by

$$(83) \quad [\xi, \eta]_R := [Y_\xi, Y_\eta](e) ,$$

and get a Lie algebra isomorphism $T_e G \cong \mathcal{X}_R(G)$.

- (B) But the two Lie brackets, eq. 73 and 83, on $T_e G$ are different. In fact one can show that:

- (i): X_ξ and Y_ξ are related by

$$(84) \quad I_* X_\xi = -Y_\xi$$

where $I : G \rightarrow G$ is the inversion map $I(g) := g^{-1}$, and I_* is the push-forward on vector fields induced by I , cf. eq. 31, i.e.

$$(85) \quad (I_* X_\xi)(g) := (T I \circ X_\xi \circ I^{-1})(g) .$$

Besides, since I is a diffeomorphism, eq. 84 makes I_* a vector space isomorphism.

- (ii): It follows from eq. 84 that

$$(86) \quad [X_\xi, X_\eta](e) = -[Y_\xi, Y_\eta](e) \quad ; \quad \text{so} \quad [\xi, \eta] = -[\xi, \eta]_R .$$

Finally, a remark about physics. In applications to physics, G is usually the group of symmetries of a physical system, and so a vector field on G is the infinitesimal generator of a one-parameter group of symmetries. For mechanics, we saw this repeatedly in Section 2, especially as regards the group of translations and rotations about the origin, in physical space \mathbb{R}^3 . This Subsection's isomorphism between the Lie algebra \mathfrak{g} and left-invariant vector fields on G means that we can think of \mathfrak{g} also as consisting of infinitesimal symmetries of the system. (The $\xi \in \mathfrak{g}$ are also called generators of the group G .)

3.4.3 Examples, subgroups and subalgebras

I begin with the first two of Section 3.4.1's three examples. That will prompt a little more theory, which will enable us to deal efficiently in the next Subsection with the third example, viz. the rotation group.

(1): *Examples*

- (i) $G := \mathbb{R}^n$ under addition. G is abelian so that left and right translation coincide. The invariant vector fields are just the constant vector fields, so that $\mathcal{X}_L(G) \equiv \mathcal{X}_R(G) \cong \mathbb{R}^n$. So the tangent space at the identity $T_e G$, i.e. the Lie algebra \mathfrak{g} , is itself \mathbb{R}^n . The bracket structure is wholly degenerate: for all invariant vector fields X, Y , $[X, Y] = 0$; and for all $\xi, \eta \in \mathfrak{g}$, $[\xi, \eta] = 0$.
- (ii) $G := GL(n, \mathbb{R})$, the general linear group. Since G is open in $End(\mathbb{R}^n, \mathbb{R}^n)$, the vector space of all linear maps on \mathbb{R}^n ('*End*' for 'endomorphism'), G 's Lie algebra, as a vector space, is $End(\mathbb{R}^n, \mathbb{R}^n)$; (cf. example (i)). To compute what the Lie bracket is, we first note that any $\xi \in End(\mathbb{R}^n, \mathbb{R}^n)$ defines a corresponding vector field on $GL(n, \mathbb{R})$ by

$$(87) \quad X_\xi : A \in GL(n, \mathbb{R}) \mapsto A\xi \in End(\mathbb{R}^n, \mathbb{R}^n) .$$

Besides, X_ξ is left-invariant, since for every $B \in GL(n, \mathbb{R})$, the left translation

$$(88) \quad L_B : A \in GL(n, \mathbb{R}) \mapsto BA \in GL(n, \mathbb{R})$$

is linear, and so

$$(89) \quad X_\xi(L_B A) = BA\xi = T_A L_B X_\xi(A) .$$

Applying now eq. 59 at the identity $I \in GL(n, \mathbb{R})$ to the definition of the bracket in the Lie algebra, eq. 73, we have:

$$(90) \quad [\xi, \eta] := [X_\xi, X_\eta](I) = \mathbf{D}X_\eta(I) \cdot X_\xi(I) - \mathbf{D}X_\xi(I) \cdot X_\eta(I).$$

But $X_\eta A = A\eta$ is linear in A , so $\mathbf{D}X_\eta(I) \cdot B = B\eta$. This means that

$$(91) \quad \mathbf{D}X_\eta(I) \cdot X_\xi(I) = \xi\eta ;$$

and similarly

$$(92) \quad \mathbf{D}X_\xi(I) \cdot X_\eta(I) = \eta\xi.$$

So the Lie algebra $\text{End}(\mathbb{R}^n, \mathbb{R}^n)$ has the usual matrix commutator as its bracket: $[\xi, \eta] = \xi\eta - \eta\xi$. This Lie algebra is often written $\mathfrak{gl}(n, \mathbb{R})$.

Let us apply to this example, result (2) from Section 3.4.2.2. In short, the result said that left-invariant vector fields correspond (by exponentiation through $e \in G$) to connected one-parameter subgroups of G . To find the one-parameter subgroup $\exp(\tau X_\xi)(e)$ of $GL(n, \mathbb{R})$, we take the matrix entries x_{ij} , ($i, j = 1, \dots, n$) as the n^2 coordinates on $GL(n, \mathbb{R})$, so that the tangent space at the identity matrix I is the set of vectors

$$(93) \quad \Sigma_{ij} \xi_{ij} \frac{\partial}{\partial x_{ij}} \Big|_I$$

with $\xi = (\xi_{ij})$ an arbitrary matrix. For given ξ , $\exp(\tau X_\xi)e$ is found by integrating the n^2 ordinary differential equations

$$(94) \quad \frac{dx_{ij}}{d\tau} = \Sigma_k \xi_{ik} x_{kj} \quad ; \quad x_{ij}(0) = \delta_{ij}.$$

The solution is just the matrix exponential:

$$(95) \quad X(\tau) = \exp(\tau\xi).$$

More generally, let us return to Section 3.4.1's idea of a matrix Lie group. For a matrix Lie group G , the definition of its Lie algebra can be given as:

$$(96) \quad \mathfrak{g} = \{ \text{the set of matrices } \xi = \phi'(0) : \phi \text{ a differentiable map } : \mathbb{R} \rightarrow G, \phi(0) = e_G \}.$$

The deduction of the structure of the Lie algebra then proceeds straightforwardly. In particular, we get the result that the one-parameter subgroup generated by $\xi \in \mathfrak{g}$ is given by matrix exponentials, as in eq. 95: the group is $\{\exp(\tau\xi) : \tau \in \mathbb{R}\}$.

This result will help us compute our third example: finding the Lie algebra of the rotation group. But for that example, it is worth first developing a little the result (2) from Section 3.4.2.2: i.e. the correspondence between left-invariant vector fields and connected one-parameter subgroups of G .

(2): *More theory:*

First, a warning remark. We will later need to take notice of the fact that a subgroup, even a one-parameter subgroup, of a Lie group G need not be a submanifold of G . Here we recall Section 3.3.1's definitions of immersion and embedding. Accordingly, we now define a subgroup H of a Lie group G to be a *Lie subgroup* of G if the inclusion map $i : H \rightarrow G$ is an injective immersion.

Just as we saw in Section 3.3.1 that not every injective immersion is an embedding, so also there are examples of Lie subgroups that are not submanifolds.

Example: the torus \mathbb{T}^2 can be made into a Lie group in a natural way (exercise: do this!); the one-parameter subgroups on the torus \mathbb{T}^2 that wind densely on the torus are Lie subgroups that are not submanifolds. (For more details about this example, cf. [Arnold, 1973, 160–167] or Arnold [1989, 72–74] or Butterfield [2004a, Section 2.1.3.B].)

But it turns out that being closed is a sufficient, and necessary, further condition. That is:

If H is a closed subgroup of a Lie group G , then H is a submanifold of G and in particular a Lie subgroup. And conversely, if H is a Lie subgroup that is also a submanifold, then H is closed.

Result (2) from Section 3.4.2.2, i.e. the correspondence between one-dimensional subgroups of G and one-dimensional subspaces (and so subalgebras) of \mathfrak{g} , generalizes to higher-dimensional subgroups and subalgebras. That is to say:

If $H \subset G$ is a Lie subgroup of G , then its Lie algebra $\mathfrak{h} := \mathfrak{g}(H)$ is a subalgebra of $\mathfrak{g} \equiv \mathfrak{g}(G)$. In fact

$$(97) \quad \mathfrak{h} = \{\xi \in \mathfrak{g} : \exp(\tau X_\xi)(e) \in H, \text{ for all } \tau \in \mathbb{R}\}.$$

And conversely, if \mathfrak{h} is any m -dimensional subalgebra of \mathfrak{g} , then there is a unique connected m -dimensional Lie subgroup H of G with Lie algebra \mathfrak{h} .

The proof of the first two statements uses result (1) of Section 3.4.2.2. For the third, i.e. converse, statement, the main idea is that \mathfrak{h} defines m vector fields on G that are linearly independent and in involution, so that one can apply Frobenius' theorem to infer an integral submanifold. One then has to prove that H is a Lie subgroup: Olver [2000, Theorem 1.51] and Marsden and Ratiu [1999, 279–280] give details and references. (Historical note: to see that this result, sometimes called Lie's 'third fundamental theorem', is close to what Lie himself called the main theorem of his theory of groups, cf. [Hawkins, 2000, 83].)

This general correspondence between Lie subgroups and Lie subalgebras prompts the question whether every finite-dimensional Lie algebra \mathfrak{g} is the Lie algebra of a Lie group. The answer is Yes. Besides, the question reduces to the case of a matrix Lie group (i.e. a Lie subgroup of $GL(n, \mathbb{R})$), in the sense that: every finite-dimensional Lie algebra \mathfrak{g} is isomorphic to a subalgebra of $\mathfrak{gl}(n, \mathbb{R})$, for some n . But be warned: this does not imply (and it is not true) that every Lie group is realizable as a matrix Lie group, i.e. that every Lie group is isomorphic to a Lie subgroup of $GL(n, \mathbb{R})$.

This general correspondence also simplifies greatly the computation of the Lie algebras of Lie groups, for example $H := SO(3)$, that are Lie subgroups of $GL(n, \mathbb{R})$. We only need to combine it with example (ii) above, that $\mathfrak{gl}(n, \mathbb{R})$ is $End(\mathbb{R}^n, \mathbb{R}^n)$ with the usual matrix commutator as its bracket: $[\xi, \eta] = \xi\eta - \eta\xi$.

Thus we infer that the Lie algebra of $SO(3)$, written $\mathfrak{so}(3)$, is a subalgebra

of $\text{End}(\mathbb{R}^n, \mathbb{R}^n)$ with the matrix commutator as bracket. Besides, we can identify $\mathfrak{so}(3)$ by looking at all the one-dimensional subgroups of G contained in it. Combining eq. 95 and 97, we have

$$(98) \quad \mathfrak{so}(3) = \{\xi \in \mathfrak{gl}(n, \mathbb{R}) : \text{the matrix exponential } \exp(\tau\xi) \in SO(3), \forall \tau \in \mathbb{R}\}.$$

With this result in hand, we can now compute $\mathfrak{so}(3)$.

3.4.4 The Lie algebra of the rotation group

Our first aim is to calculate the Lie algebra $\mathfrak{so}(3)$ (also written: $so(3)$) of $H := SO(3)$, the rotation group. This will lead us back to Section 3.2.1.1 correspondence between anti-symmetric matrices and vectors in \mathbb{R}^3 .

$SO(3)$ is represented by 3×3 orthogonal matrices of determinant 1. So the requirement in eq. 98 becomes, now writing e , not \exp :

$$(99) \quad (e^{\tau\xi})(e^{\tau\xi})^T = I \quad \text{and} \quad \det(e^{\tau\xi}) = 1.$$

Differentiating the first equation with respect to τ and setting $\tau = 0$ yields

$$(100) \quad \xi + \xi^T = 0.$$

So ξ must be anti-symmetric, i.e. represented by an anti-symmetric matrix. Conversely, for any such anti-symmetric matrix ξ , we can show that $\det(e^{\tau\xi}) = 1$. So, indeed:

$$(101) \quad \mathfrak{so}(3) = \{3 \times 3 \text{ antisymmetric matrices}\}.$$

Notice that the argument is independent of choosing $n = 3$. It similarly computes $\mathfrak{so}(n)$ for any integer n :

$$(102) \quad \mathfrak{so}(n) = \{n \times n \text{ antisymmetric matrices}\}.$$

Thus the rotations on euclidean space \mathbb{R}^n of any dimension n are generated by the Lie algebra of $n \times n$ anti-symmetric matrices.

This justifies our assertion at the end of Section 3.2.1.1 that the rotation group in three dimensions is special in being representable by vectors in the space on which it acts, i.e. \mathbb{R}^3 . For as we have just seen, in general the infinitesimal generators of rotations are anti-symmetric matrices, which in n dimensions have $n(n-1)/2$ independent components. But only for $n = 3$ does this equal n .

Remark: An informal computation of $\mathfrak{so}(3)$, based on the idea that higher-order terms in $e^{\tau\xi}$ can be neglected (cf. the physical idea that ξ represents an infinitesimal rotation), goes as follows.

For $(I + \tau\xi)$ to be a rotation requires that

$$(103) \quad (I + \tau\xi)(I + \tau\xi)^T = I \quad \text{and} \quad \det(I + \xi\tau) = 1.$$

Dropping higher-order terms, the first equation yields

$$(104) \quad I + \tau(\xi + \xi^T) = I \quad \text{i.e.} \quad \xi + \xi^T = 0.$$

Besides, the second equation in eq. 103 yields no further constraint, since for any anti-symmetric matrix ξ written as (cf. eq. 47)

$$(105) \quad \xi = \begin{pmatrix} 0 & -\xi_3 & \xi_2 \\ \xi_3 & 0 & -\xi_1 \\ -\xi_2 & \xi_1 & 0 \end{pmatrix},$$

we immediately compute that $\det(I + \xi\tau) = 1 + \tau^2(\xi_1^2 + \xi_2^2 + \xi_3^2)$. So, dropping higher-order terms, $\det(I + \xi\tau) = 1$. In short, we again conclude that

$$(106) \quad \mathfrak{so}(3) = \{3 \times 3 \text{ antisymmetric matrices}\}.$$

For later use (e.g. Sections 4.4 and 4.5.1), we note that the three matrices

$$(107) \quad A^x = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \quad A^y = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad A^z = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

span $\mathfrak{so}(3)$, and generate the one-parameter subgroups

$$(108) \quad R_\theta^x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}, \quad R_\theta^y = \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix},$$

$$R_\theta^z = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

representing anticlockwise rotation around the respective coordinate axes in the physical space \mathbb{R}^3 .

Having computed $\mathfrak{so}(3)$ to consist of antisymmetric matrices, we can use Section 3.2.1.1's correspondence between these and vectors in \mathbb{R}^3 so as to realize $\mathfrak{so}(3)$ as vectors with the Lie bracket as vector multiplication. With these realizations in hand, we can readily obtain several further results about rotations. We will not need any. But a good example, which uses eq. 48's isomorphism Θ from vectors $\omega \in \mathbb{R}^3$ to matrices $A \in \mathfrak{so}(3)$, is as follows: —

$\exp(\tau \Theta(\omega))$ is a rotation about the axis ω by the angle $\tau \|\omega\|$.

We can now begin to see the point of this Chapter's second motto (from Arnold), that the elementary theory of the rigid body confuses six conceptually different three-dimensional spaces. For our discussion has already distinguished three of the six spaces which Arnold lists (in a different notation). Namely, we have just distinguished:

- (i) \mathbb{R}^3 , especially when taken as physical space; from (ii) $\mathfrak{so}(3) \equiv T_e(SO(3))$, the generators of rotations; though they are isomorphic as Lie algebras, by eq. 48's bijection Θ from vectors $\omega \in \mathbb{R}^3$ to matrices $A \in \mathfrak{so}(3)$;

- (ii) $\mathfrak{so}(3) \equiv T_e(SO(3))$ from its isomorphic copy under the derivative of left translation by g (i.e. under $(L_g)_*$), viz. $T_g(SO(3))$: cf. eq. 69. (In the motto, Arnold writes g for $\mathfrak{so}(3)$ and G for $SO(3)$.)

In Section 5.2.4 we will grasp (even without developing the theory of the rigid body!) the rest of the motto. That is, we will see why Arnold also mentions the three corresponding dual spaces, \mathbb{R}^{3*} , $\mathfrak{so}(3)^*$ and $T_g^*(SO(3))$. But we can already say more about the two tangent spaces $\mathfrak{so}(3) \equiv T_e(SO(3))$ and $T_g(SO(3))$, in connection with the idea that for a pivoted rigid body, the configuration space can be taken as $SO(3)$; (cf. (3) of Section 2.2). We will show that there are two isomorphisms from $T_g(SO(3))$ to $T_e(SO(3))$ that are natural, not only in the mathematical sense of being basis-independent but also in the sense of having a physical interpretation. Namely, they represent the computation of the angular velocity from the Lagrangian generalized velocity, i.e. \dot{q} . In effect, one isomorphism computes the angular velocity's components with respect to an orthonormal frame fixed in space (called *spatial coordinates*); and the other computes it with respect to a frame fixed in the rigid body (*body coordinates*). In fact, these isomorphisms are the derivatives of right and left translation, respectively; (cf. eq. 67 and 68).

So suppose a pivoted rigid body has a right-handed orthonormal frame $\{a, b, c\}$ fixed in it. We can think of the three unit vectors a, b, c as column vectors in \mathbb{R}^3 . Arranging them in a 3×3 matrix $g := (a \ b \ c) \in GL(3, \mathbb{R})$, we get a matrix that maps the unit x -vector e_1 to a , the unit y -vector e_2 to b , etc. That is: g maps the standard frame e_1, e_2, e_3 to a, b, c , and g is an orthogonal matrix: $g \in SO(3) = \{g \in GL(3, \mathbb{R}) \mid \tilde{g}g = I\}$. Thus g represents the configuration of the body, and the configuration space is $SO(3)$.

By differentiating the condition $\tilde{g}g = I$, we deduce that the tangent space at a specific g , $T_g(SO(3))$, i.e. the space of velocities \dot{g} , is the 3-dimensional vector subspace of $GL(3, \mathbb{R})$:

$$(109) \quad T_g(SO(3)) = \{\dot{g} \in GL(3, \mathbb{R}) \mid \dot{\tilde{g}}g + \tilde{g}\dot{g} = 0\}$$

Now recall examples (ii) and (iii) of Section 3.2.1.1. We saw there that though the angular velocity of the body is usually taken to be the vector ω such that, with our "body-vectors" a, b, c ,

$$(110) \quad \dot{a} = \omega \wedge a, \quad \dot{b} = \omega \wedge b, \quad \dot{c} = \omega \wedge c :$$

we can instead encode the angular velocity by the antisymmetric matrix $A := \Theta(\omega) \in \mathfrak{g} \equiv T_e(SO(3))$. As we saw, eq. 110 then becomes

$$(111) \quad \dot{a} = \Theta(\omega)a, \quad \dot{b} = \Theta(\omega)b, \quad \dot{c} = \Theta(\omega)c :$$

or equivalently the matrix equation for the configuration $g = (a \ b \ c)$,

$$(112) \quad \dot{g} \equiv (\dot{a} \ \dot{b} \ \dot{c}) = \Theta(\omega)g ; \quad \text{i.e.} \quad \Theta(\omega) = \dot{g}g^{-1} .$$

Thus we see that the map from $T_g(SO(3))$ to $\mathfrak{g} = T_e(SO(3))$

$$(113) \quad \dot{g} \in T_g(SO(3)) \mapsto \dot{g}g^{-1} \equiv \dot{g}\tilde{g} \in \mathfrak{g}$$

maps the generalized velocity \dot{g} to the angular velocity $\Theta(\omega)$. This is the angular velocity represented in the usual elementary way, with respect to coordinates fixed in space. One immediately checks that it is an isomorphism (exercise!).

On the other hand, let us consider $\Theta(\omega)$ as a linear transformation $\Theta(\omega) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, and express it in the *body* coordinates a, b, c . This gives $g^{-1}\Theta(\omega)g \equiv g^{-1}\dot{g}$. Thus the map

$$(114) \quad \dot{g} \in T_g(SO(3)) \mapsto g^{-1}\dot{g} \equiv \tilde{g}\dot{g} \in \mathfrak{g}$$

maps the generalized velocity \dot{g} to the angular velocity expressed in body coordinates. It also is clearly an isomorphism.

Summing up: we have two natural isomorphisms that compute the angular velocity, in spatial and body coordinates respectively, from the generalized velocity \dot{g} .

Incidentally, one can verify directly that the images $\dot{g}\tilde{g}$ and $\tilde{g}\dot{g}$ of the isomorphisms eq. 113 and 114 lie in \mathfrak{g} , i.e. are antisymmetric matrices. Thus with \cdot for the elementary dot-product, we have:

$$(115) \quad g^{-1}\dot{g} \equiv \tilde{g}\dot{g} = \begin{pmatrix} \tilde{a} \\ \tilde{b} \\ \tilde{c} \end{pmatrix} (\dot{a} \ \dot{b} \ \dot{c}) = \begin{pmatrix} 0 & a \cdot \dot{b} & a \cdot \dot{c} \\ b \cdot \dot{a} & 0 & b \cdot \dot{c} \\ c \cdot \dot{a} & c \cdot \dot{b} & 0 \end{pmatrix} .$$

This is an antisymmetric matrix, since differentiating $a \cdot b = b \cdot c = a \cdot c = 0$ with respect to time gives $a \cdot \dot{b} + \dot{a} \cdot b = 0$ etc. Finally, we deduce that $\dot{g}\tilde{g}$ is antisymmetric from the facts that $\dot{g}\tilde{g} = g(g^{-1}\dot{g})g^{-1}$ and antisymmetry is preserved by conjugation by g .

We end this Subsection with two incidental remarks; (they will not be used in what follows).

- (1) In Section 2.1.1, we could have specialized the discussion from a symplectic manifold to a symplectic vector space, i.e. a (real, finite-dimensional) vector space equipped with a non-degenerate anti-symmetric bilinear form $\omega : Z \times Z \rightarrow \mathbb{R}$. It follows that Z is of even dimension. The question then arises which linear maps $A : Z \rightarrow Z$ preserve the normal form of ω given by eq. 4. It is straightforward to show that this is equivalent to A preserving the form of Hamilton's equations (for any Hamiltonian); so that these maps A are called *symplectic* (or *canonical*, or *Poisson*). The set of all such maps form a Lie group, the *symplectic group*, written $\text{Sp}(Z, \omega)$. But since this Chapter will not need the theory of canonical transformations, I leave the study of $\text{Sp}(Z, \omega)$'s structure as an exercise! (For details, cf. e.g. Abraham and Marsden [1978, 167–174], Marsden and Ratiu [1999, : 69–72, 293–299].)
- (2) Finally, a glimpse of the infinite-dimensional manifolds that this Chapter has foresworn. Consider the infinite-dimensional Lie group $\text{Diff}(M)$ of all diffeomorphisms on M . An element of its Lie algebra, i.e. a vector $A \in T_e(\text{Diff}(M))$, is a vector field, or equivalently a flow, on M . Besides, the

Lie bracket in this Lie algebra $T_e(\text{Diff}(M))$, as defined by eq. 73 turns out to be the usual Lie bracket of the vector fields on M , as defined in Section 3.2.2.

4 ACTIONS OF LIE GROUPS

We turn to actions of Lie groups on manifolds. The notions, results and examples in this Section will be crucial from Section 5.4 onwards. Fortunately, the foregoing provides several examples of the notions and results we need. Section 4.1 will give basic material, including the crucial notion of cotangent lifts. Sections 4.2 and 4.3 describe conditions for orbits and quotient spaces to be manifolds. Section 4.4 describes actions infinitesimally, i.e. in terms of their infinitesimal generators. Section 4.5 presents two important representations of a Lie group, its adjoint and co-adjoint representations, on its Lie algebra \mathfrak{g} and on the dual \mathfrak{g}^* respectively. Finally, Section 4.6 gathers some threads concerning our central, recurring example, viz. the rotation group.

4.1 Basic definitions and examples

A *left action* of a Lie group G on a manifold M is a smooth map $\Phi : G \times M \rightarrow M$ such that:

- (i) $\Phi(e, x) = x$ for all $x \in M$
- (ii) $\Phi(g, \Phi(h, x)) = \Phi(gh, x)$ for all $g, h \in G$ and all $x \in M$.

We sometimes write $g \cdot x$ for $\Phi(g, x)$.

Similarly, a *right action* of a Lie group G on a manifold M is a smooth map $\Psi : M \times G \rightarrow M$ satisfying (i) $\Psi(x, e) = x$ and (ii) $\Psi(\Psi(x, g), h) = \Psi(x, gh)$. We sometimes write $x \cdot g$ for $\Psi(x, g)$.

It is convenient to also have a subscript notation. For every $g \in G$, we define

$$(116) \quad \Phi_g : M \rightarrow M \quad : \quad x \mapsto \Phi(g, x).$$

In this notation, (i) becomes $\Phi_e = id_M$ and (ii) becomes $\Phi_{gh} = \Phi_g \circ \Phi_h$. For right actions, (ii) becomes $\Psi_{gh} = \Psi_h \circ \Psi_g$.

One immediately verifies that any left action Φ of G on a manifold M , $g \mapsto \Phi_g : M \rightarrow M$, defines a right action Ψ by

$$(117) \quad g \mapsto \Psi_g := \Phi_{g^{-1}} : M \rightarrow M \quad ; \quad \text{i.e.} \quad \Psi : (x, g) \in M \times G \mapsto \Phi(g^{-1}, x) \in M.$$

(Use the fact that in G , $(gh)^{-1} = h^{-1}g^{-1}$.) Similarly, a right action defines a left action, by taking the inverse in G . We will occasionally make use of this left-right “flip”.

The definition of left action is equivalent to saying that the map $g \mapsto \Phi_g$ is a homomorphism of G into $\text{Diff}(M)$, the group of diffeomorphisms of M . In the

special case where M is a Banach space V and each $\Phi_g : V \rightarrow V$ is a continuous linear transformation, the action of G on V is called a *representation* of G on V .

The *orbit* of $x \in M$ (under the action Φ) is the set

$$(118) \quad \text{Orb}(x) = \{\Phi_g(x) : g \in G\} \subset M.$$

The action is called *transitive* if there is just one orbit, i.e. for all $x, y \in M$, there is a $g \in G$ such that $g \cdot x = y$. It is called *effective* (or *faithful*) if $\Phi_g = \text{id}_M$ implies $g = e$, i.e. if $g \mapsto \Phi_g$ is one-to-one. It is called *free* if it has no fixed points for any $g \neq e$: that is, $\Phi_g(x) = x$ implies $g = e$. In other words, it is free if for each $x \in M$, $g \mapsto \Phi_g(x)$ is one-to-one. (So: every free action is faithful.)

4.1.0.1 Examples; cotangent lifts We begin with geometric examples; and then return to mechanics, giving first some general theory, followed by some examples.

(2): *Geometric examples:*

- (i) $SO(3)$ acts on \mathbb{R}^3 by $(A, x) \mapsto Ax$. The action is faithful. But it is neither free (each rotation fixes the points on its axis) nor transitive (the orbits are the spheres centred at the origin).
- (ii) $GL(n, \mathbb{R})$ acts on \mathbb{R}^n by $(A, x) \mapsto Ax$. The action is faithful, not free, and “almost transitive”: the zero subspace $\{0\}$ is an orbit, and so is $\mathbb{R}^n - \{0\}$.
- (iii) Suppose X is a vector field on M which is *complete* in the sense that the flow $\phi_X(\tau)$ of eq. 35 is defined for all $\tau \in \mathbb{R}$. Then this flow defines an action of \mathbb{R} on M .

We turn to two examples which will be central, and recurring, in our discussion of symplectic reduction.

- (iv) Left translation by each $g \in G$, $L_g : h \in G \mapsto gh \in G$ (cf. eq. 67), defines a left action of G on itself. Since G is a group, it is transitive and free (and so faithful). Similarly, right translation, $g \mapsto R_g$ with $R_g : h \in G \mapsto hg \in G$, defines a *right* action. And $g \mapsto R_{g^{-1}}$ defines a *left* action; cf. eq. 117. One readily proves that left translation lifts to the tangent bundle TG as a left action. That is: one verifies by the chain rule that

$$(119) \quad \Phi_g : TG \rightarrow TG \quad : \quad v \equiv v_h \in T_h G \mapsto (T_h L_g)(v) \in T_{gh} G$$

defines a left action on TG . Similarly, right translation lifts to a right action on TG . But our interest in Hamiltonian mechanics of course makes us more interested in *cotangent* lifts. See (2) below for the general definitions, and example (viii) in (3) below for the cotangent lift of left translation.

- (v) G acts on itself by *conjugation* (inner automorphism): $g \mapsto K_g := R_{g^{-1}} \circ L_g$. That is: $K_g : h \in G \mapsto ghg^{-1} \in G$. Each K_g is an isomorphism of G . The

orbits are conjugacy classes. Section 4.5 will introduce two “differentiated versions” of action by conjugation, viz. the adjoint and co-adjoint actions, which will be important in symplectic reduction.

(2): *Hamiltonian symmetries and cotangent lifts:*

We turn to Hamiltonian mechanics. Following the discussion in Section 2.1.3, we say: given a Hamilton system (M, ω, H) with (M, ω) a symplectic manifold and $H : M \rightarrow \mathbb{R}$, a *Hamiltonian group of symmetries* is a Lie group G acting on M such that each $\Phi_g : M \rightarrow M$ preserves both ω and H . Then the simplest possible examples are spatial translations and-or rotations acting on the free particle. The details of these examples, (vi) and (vii) below, will be clearer if we first develop some general theory.

This theory will illustrate the interaction between the left-right contrast for actions, and the tangent-cotangent contrast for bundles. Besides, both the general theory and the examples’ details will carry over straightforwardly, i.e. component by component, to the case of N particles interacting by Newtonian gravity, discussed in Section 2.3.2: the action defined on a single particle is just repeated for each of the N particles.

So we will take $M := (\mathbb{R}^3) \times (\mathbb{R}^3)^*$, $\omega := dq^i \wedge dp^i$, $H := p^2/2m$. In the first place, both translations (by $\mathbf{x} \in \mathbb{R}^3$) and rotations (by $A \in SO(3)$) act on the configuration space $Q = \mathbb{R}^3$. We have actions of \mathbb{R}^3 and $SO(3)$ on \mathbb{R}^3 by

$$(120) \quad \Phi_{\mathbf{x}}(\mathbf{q}) = \mathbf{q} + \mathbf{x} \quad ; \quad \Phi_A(\mathbf{q}) = A\mathbf{q} .$$

But these actions *lift* to the cotangent bundle $T^*Q = (\mathbb{R}^3) \times (\mathbb{R}^3)^* \cong \mathbb{R}^6$; (as mentioned in Section 2.3.2). The lift of these actions is defined using a result that does not use the notion of an action. Namely:

Any diffeomorphism $f : Q_1 \rightarrow Q_2$ induces a *cotangent lift* $T^*f : T^*Q_2 \rightarrow T^*Q_1$ (i.e. in the opposite direction) which is symplectic, i.e. maps the canonical one-form, and so symplectic form, on T^*Q_2 to that of T^*Q_1 .

To define the lift of an action, it is worth going into detail about the definition of T^*f . (But I will not prove the result just stated; for details, cf. [Marsden and Ratiu, 1999, Section 6.3].)

The idea is that T^*f is to be the “pointwise adjoint” of the tangent map $Tf : TQ_1 \rightarrow TQ_2$ (eq. 29). That is: we define T^*f in terms of the contraction of its value, for an arbitrary argument $\alpha \in T_{q_2}^*Q_2$, with an arbitrary tangent vector $v \in T_{f^{-1}(q_2)}Q_1$. (Here it will be harmless to (follow many presentations and) conflate a point in T^*Q_2 , i.e. strictly speaking a pair (q_2, α) , $q_2 \in Q_2, \alpha \in T_{q_2}^*Q_2$, with its form α . And similarly it will be harmless to conflate a point (q_1, v) in TQ_1 with its vector $v \in T_{q_1}Q_1$.)

We recall that any finite-dimensional vector space is naturally, i.e. basis-independently, isomorphic to its double dual: $(V^*)^* \cong V$; and we will use angle brackets $\langle ; \rangle$ for the natural pairing between V and V^* . So we define $T^*f : T^*Q_2 \rightarrow T^*Q_1$ by requiring:

$$(121) \quad \langle (T^*f)(\alpha); v \rangle := \langle \alpha; (Tf)(v) \rangle, \quad \forall \alpha \in T_{q_2}^*Q_2, v \in T_{f^{-1}(q_2)}Q_1 .$$

NB: Because T^*f “goes in the opposite direction”, the composition of lift with function-composition involves a reversal of the order. That is: if $Q_1 = Q_2 \equiv Q$ and f, g are two diffeomorphisms of Q , then

$$(122) \quad T^*(f \circ g) = T^*g \circ T^*f .$$

With this definition of T^*f , a left action Φ of G on the manifold Q induces for each $g \in G$ the cotangent lift of $\Phi_g : Q \rightarrow Q$. That is: we have the map

$$(123) \quad T^*\Phi_g \equiv T^*(\Phi_g) : T^*Q \rightarrow T^*Q, \quad \text{with } \alpha \in T_q^*Q \mapsto (T^*\Phi_g)(\alpha) \in T_{g^{-1} \cdot q}^*Q .$$

Now consider the map assigning to each $g \in G$, $T^*\Phi_g$:

$$(124) \quad g \in G \mapsto T^*\Phi_g : T^*Q \rightarrow T^*Q .$$

To check that this is indeed an action of G on T^*Q , we first check that since $\Phi_e = id_Q$, $T\Phi_e : TQ \rightarrow TQ$ is id_{TQ} and $T^*(\Phi_e)$ is id_{T^*Q} . But beware: eq. 122 yields

$$(125) \quad T^*\Phi_{gh} = T^*(\Phi_g \circ \Phi_h) = T^*\Phi_h \circ T^*\Phi_g ,$$

so that eq. 124 defines a *right* action.

But here we recall that any left action defines a right action by using the inverse; cf. eq. 117. Combining this with the idea of the cotangent lift of an action on Q , we get:

The left action Φ on Q defines, not only the right action eq. 124 on T^*Q , but also a *left* action on T^*Q , viz. by

$$(126) \quad g \in G \mapsto \Psi_g := T^*(\Phi_{g^{-1}}) : T^*Q \rightarrow T^*Q .$$

For since $(gh)^{-1} = h^{-1}g^{-1}$,

$$(127) \quad \Psi_{gh} \equiv T^*(\Phi_{(gh)^{-1}}) = T^*(\Phi_{h^{-1}g^{-1}}) = T^*(\Phi_{h^{-1}} \circ \Phi_{g^{-1}}) = T^*\Phi_{g^{-1}} \circ T^*\Phi_{h^{-1}} \equiv \Psi_g \circ \Psi_h .$$

In short, the two reversals of order cancel out. This sort of left-right flip will recur in some important contexts in the following, in particular in Sections 6.5 and 7.

(3): *Mechanical examples:*

So much by way of generalities. Now we apply them to translations and rotations of a free particle, to rotations of a pivoted rigid body, and to N point-particles.

(vi): Let the translation group $G = (\mathbb{R}^3, +)$ act on the free particle’s configuration space $Q = \mathbb{R}^3$ by

$$(128) \quad \Phi_{\mathbf{x}}(\mathbf{q}) = \mathbf{q} + \mathbf{x} .$$

Since G is abelian, the distinction between left and right actions of G collapses. (And if we identify G with Q , this is left=right translation by \mathbb{R}^3 on itself, i.e. example (iv) again: and so transitive and free.) But of course the lifted actions we have defined, “with g ” and “with g^{-1} ”, eq. 124 and 126 respectively, remain distinct actions.

Then, writing $\alpha = (\mathbf{q}, \mathbf{p}) \in T_{\mathbf{q}}^*Q$, and using the fact that $T\Phi_{\mathbf{x}}(\mathbf{q} - \mathbf{x}, \dot{\mathbf{q}}) = (\mathbf{q}, \dot{\mathbf{q}})$, we see that eq. 121 implies that: first,

$$(129) \quad T^*(\Phi_{\mathbf{x}})(\mathbf{q}, \mathbf{p}) \in T_{\mathbf{q}-\mathbf{x}}^*Q ;$$

and second, that for all $\dot{\mathbf{q}} \in T_{\mathbf{q}-\mathbf{x}}Q$,

$$(130) \quad \langle T^*(\Phi_{\mathbf{x}})(\mathbf{q}, \mathbf{p}); (\mathbf{q} - \mathbf{x}, \dot{\mathbf{q}}) \rangle = \langle (\mathbf{q}, \mathbf{p}); (\mathbf{q}, \dot{\mathbf{q}}) \rangle \equiv \mathbf{p}(\dot{\mathbf{q}}) .$$

For eq. 130 to hold for all $\dot{\mathbf{q}} \in T_{\mathbf{q}-\mathbf{x}}Q$ requires that $T^*(\Phi_{\mathbf{x}})(\mathbf{q}, \mathbf{p})$ does not affect \mathbf{p} , i.e.

$$(131) \quad T^*(\Phi_{\mathbf{x}})(\mathbf{q}, \mathbf{p}) = (\mathbf{q} - \mathbf{x}, \mathbf{p}) .$$

So this is the lifted action “with g ”, corresponding to eq. 124. Similarly, the lifted action “with g^{-1} ”, corresponding to eq. 126, is: $\Psi_{\mathbf{x}}(\mathbf{q}, \mathbf{p}) := T^*(\Phi_{-\mathbf{x}})(\mathbf{q}, \mathbf{p}) = (\mathbf{q} + \mathbf{x}, \mathbf{p})$.

One readily checks that these lifted actions preserve both $\omega = dq^i \wedge dp^i$ (an exercise in manipulating the exterior derivative) and $H := p^2/2m$. So we have a Hamiltonian symmetry group. The action is not transitive: the orbits are labelled by their values of $\mathbf{p} \in (\mathbb{R}^3)^*$. But it is free.

(vii): Let $SO(3)$ act on the left on $Q = \mathbb{R}^3$ by

$$(132) \quad \Phi_A(\mathbf{q}) = A\mathbf{q} .$$

(This is example (i) again.) Let us lift this action “with g ”, i.e. eq. 124, so as to get a right action on T^*Q .

As in example (vi), we write $\alpha = (\mathbf{q}, \mathbf{p}) \in T_{\mathbf{q}}^*Q$. Using the fact that $T\Phi_A(\mathbf{q}, \dot{\mathbf{q}}) = (A\mathbf{q}, A\dot{\mathbf{q}})$, eq. 121 then implies that: first,

$$(133) \quad T^*(\Phi_A)(\mathbf{q}, \mathbf{p}) \in T_{A^{-1}\mathbf{q}}^*Q ;$$

and second, that for all $\dot{\mathbf{q}} \in T_{A^{-1}\mathbf{q}}Q$,

$$(134) \quad \langle T^*(\Phi_A)(\mathbf{q}, \mathbf{p}); (A^{-1}\mathbf{q}, \dot{\mathbf{q}}) \rangle = \langle (\mathbf{q}, \mathbf{p}); (\mathbf{q}, A\dot{\mathbf{q}}) \rangle \equiv \mathbf{p}(A\dot{\mathbf{q}}) \equiv p_i A_j^i \dot{q}^j .$$

For eq. 134 to hold for all $\dot{\mathbf{q}} \in T_{A^{-1}\mathbf{q}}Q$ requires that

$$(135) \quad T^*(\Phi_A)(\mathbf{q}, \mathbf{p}) = (A^{-1}\mathbf{q}, \mathbf{p}A) ,$$

where $\mathbf{p}A$ is a row-vector. Or if one thinks of the \mathbf{p} components as a column vector, it requires:

$$(136) \quad T^*(\Phi_A)(\mathbf{q}, \mathbf{p}) = (A^{-1}\mathbf{q}, \tilde{A}\mathbf{p}) = (A^{-1}\mathbf{q}, A^{-1}\mathbf{p}) ,$$

where $\tilde{}$ represents the transpose of a matrix, and the last equation holds because A is an orthogonal matrix.

So this is the lifted action “with g ”, corresponding to eq. 124. Similarly, the lifted action “with g^{-1} ”, corresponding to eq. 126, is: $\Psi_A(\mathbf{q}, \mathbf{p}) := T^*(\Phi_{A^{-1}})(\mathbf{q}, \mathbf{p}) = (A\mathbf{q}, A\mathbf{p})$.

Again, one readily checks that these lifted actions preserve both $\omega = dq^i \wedge dp^i$ (another exercise in manipulating the exterior derivative!) and $H := p^2/2m$. So $SO(3)$ is a Hamiltonian symmetry group.

Like the original action of $SO(3)$ on Q , these actions are faithful. But they are not transitive: the orbits are labelled by the radii of two spheres centred at the origins of \mathbb{R}^3 and $(\mathbb{R}^3)^*$. And they are not free: suppose \mathbf{q} and \mathbf{p} are parallel and on the axis of rotation of A .

- (viii): Now we consider the pivoted rigid body. But unlike examples (vi) and (vii), we will consider only kinematics, not dynamics: even for a free body. That is, we will say nothing about the definitions of, and invariance of, ω and H ; for details of these, cf. e.g. [Abraham and Marsden, 1978, Sections 4.4 and 4.6] and the other references given in (3) of Section 2.2. We will in any case consider the dynamics of this example in more general terms (using momentum maps) in Sections 6.5.3 and 7.

We recall from the discussion at the end of Section 3.4.4 that the configuration space of the pivoted rigid body is $SO(3) =: G$. We also saw there that the space and body representations of the angular velocity $v = \dot{g} \in T_g G$ are given by right and left translation. Thus eq. 113 and 114 give:

$$(137) \quad v^S \equiv \dot{g}^S := T_g R_{g^{-1}}(\dot{g}) \quad \text{and} \quad v^B \equiv \dot{g}^B := T_g L_{g^{-1}}(\dot{g}) .$$

But we are now concerned with the *cotangent* lift of left (or right) translation. So let $SO(3)$ act on itself by left translation: $\Phi_g h \equiv L_g h = gh$. Let us lift this action “with g ”, i.e. eq. 124, to get a right action on T^*G . So let $\alpha \in T_h^*G$ and $(TL_g)(h, \dot{h}) = (gh, g\dot{h})$. Then eq. 121 implies that: first

$$(138) \quad (T^*L_g)(\alpha) \in T_{g^{-1}h}^*G ,$$

and second that for all $v \in T_{g^{-1}h}G$

$$(139) \quad \langle T^*(L_g)(\alpha); v \rangle = \langle \alpha; gv \rangle .$$

In other words, on analogy with eq. 131 and 135: for eq. 139 to hold for all $v \in T_{g^{-1}h}G$ requires that with $gv \in T_hG$:

$$(140) \quad T^*(L_g)(\alpha) : v \in T_{g^{-1}h}G \mapsto \alpha(gv) .$$

Similarly, the lifted action “with g^{-1} ” corresponding to eq. 126, i.e. the left action on T^*G , is

$$(141) \quad \langle T^*(L_{g^{-1}})(\alpha); v \rangle = \langle \alpha; g^{-1}v \rangle, \forall \alpha \in T_h^*G, v \in T_{gh}G$$

We will continue this example in Section 4.6, after developing more of the theory of Lie group actions.

Finally, let us sketch another mechanical example: the case of N particles with configuration space $Q := \mathbb{R}^{3N}$ interacting by Newtonian gravity — discussed in Section 2.3.2. This will combine and generalize examples (vi) and (vii); and lead on to the next Sections' discussions of orbits and quotients.

- (ix): As I mentioned above (before eq. 120), the cotangent-lifted actions of translations and rotations on a single particle carry over straightforwardly to the case of N particles: the action defined on a single particle is just repeated, component by component, for each of the N particles to give an action on $T^*Q \cong \mathbb{R}^{3N} \times (\mathbb{R}^{3N})^*$.

Furthermore, the groups of translations and rotations are subgroups of a single group, the Euclidean group E . I shall not define E exactly. Here, let it suffice to say that:

- (a) E 's component-wise action on the configuration space $Q := \mathbb{R}^{3N}$ has a cotangent lift, which is of course also component by component.
- (b) E 's cotangent-lifted action is not transitive, nor free; but it is faithful.
- (c) If we take as the Hamiltonian function the H of eq. 25, describing the particles as interacting by Newtonian gravity, then E is a Hamiltonian symmetry group. In fact, the kinetic and potential energies are separately invariant, essentially because the particles' interaction depends only on the inter-particle distances, not on their positions or orientations; cf. the discussion in Section 2.3.2.

A final comment about example (ix), which points towards the following Sections:

Recall that in Sections 2.3.3 and 2.3.4, we used this example as a springboard to discussing Relationist and Reductionist procedures, which quotiented the configuration space or phase space. But in order for the quotient spaces (and orbits) to be manifolds, and in particular for dimensions to add or subtract in a simple way, we needed to excise two classes of “special” points, before quotienting. These were: the class of symmetric configurations or states (i.e. those fixed by some element of E), and the class of collision configurations or states. For the quotienting of phase space advocated by Reductionism, the classes of states were $\delta \subset T^*\mathbb{R}^{3N}$ and $\Delta \subset T^*\mathbb{R}^{3N}$; (cf. Section 2.3.4 for definitions.)

With examples (vi) to (ix) in hand, we can now see that:

- (a) δ and Δ are each closed under the cotangent-lifted action of E on $T^*\mathbb{R}^{3N}$; i.e., each is a union of orbits. So E acts on $M := T^*\mathbb{R}^{3N} - (\delta \cup \Delta)$.

(b) E acts freely on M .

We will see in the sequel (especially in Sections 4.3.B and 5.5) that an action being free is one half (one conjunct) of an important sufficient condition for orbits and quotient spaces to be manifolds. The other conjunct will be the notion of an action being *proper*: which we will define in Section 4.3.

4.2 Quotient structures from group actions

In finite dimensions, any orbit $\text{Orb}(x)$ is an immersed submanifold of M . This can be proved directly (Abraham and Marsden [1978,: Ex. 1.6F(b), p. 51, and 4.1.22 p. 265]). But for our purposes, this is best seen as a corollary of some conditions under which quotient structures are manifolds; as follows.

The relation, $x \cong y$ if there is a $g \in G$ such that $g \cdot x = y$, is an equivalence relation, with the orbits as equivalence classes. We denote the quotient space, i.e. the set of orbits, by M/G (sometimes called the *orbit space*). We write the canonical projection as

$$(142) \quad \pi : M \rightarrow M/G, \quad x \mapsto \text{Orb}(x);$$

and we give M/G the quotient topology by defining $U \subset M/G$ to be open iff $\pi^{-1}(U)$ is open in M .

Simple examples (e.g. (ii) of Section 4.1.0.1) show that this quotient topology need not be Hausdorff. However, it is easy to show that if the set

$$(143) \quad R := \{(x, \Phi_g x) \in M \times M : (g, x) \in G \times M\}$$

is a closed subset of $M \times M$, then the quotient topology on M/G is Hausdorff.

But to ensure that M/G has a manifold structure, further conditions are required. The main one (and a much harder theorem) is:

R is a closed submanifold of $M \times M$ iff M/G is a manifold with $\pi : M \rightarrow M/G$ a submersion.

This theorem has two Corollaries which are important for us.

- (1) A map $h : M/G \rightarrow N$, from the manifold M/G , for which $\pi : M \rightarrow M/G$ is a submersion, to the manifold N , is smooth iff $h \circ \pi : M \rightarrow N$ is smooth.

This corollary has a useful implication, called *passage to the quotients*, about the notion of *equivariance*—which will be important in symplectic reduction.

A smooth map $f : M \rightarrow N$ is called *equivariant* if it respects the action of a Lie group G on the manifolds. That is: Let G act on M and N by $\Phi_g : M \rightarrow M$ and $\Psi_g : N \rightarrow N$ respectively. $f : M \rightarrow N$ is called equivariant with respect to these actions if for all $g \in G$

$$(144) \quad f \circ \Phi_g = \Psi_g \circ f.$$

That is, f is equivariant iff for all g , the following diagram commutes:

$$(145) \quad \begin{array}{ccc} M & \xrightarrow{f} & N \\ \uparrow \Phi_g & & \uparrow \Psi_g \\ M & \xrightarrow{f} & N \end{array}$$

Equivariance immediately implies that f naturally induces a map, \hat{f} say, on the quotients. That is: the map

$$(146) \quad \hat{f} : \text{Orb}(x) \in M/G \mapsto \text{Orb}(f(x)) \in N/G$$

is well-defined, i.e. independent of the chosen representative x for the orbit.

Applying the corollary we have: If $f : M \rightarrow N$ is equivariant, and the quotients M/G and M/N are manifolds with the canonical projections both submersions, then f being smooth implies that \hat{f} is smooth. This is called *passage to the quotients*.

- (2) Let H be a closed subgroup of the Lie group G . (By (2) of Section 3.4.3, this is equivalent to H being a subgroup that is a submanifold of G .) Let H act on G by left translation: $(h, g) \in H \times G \mapsto hg \in G$, so that the orbits are the right cosets Hg . Then G/H is a manifold and $\pi : G \rightarrow G/H$ is a submersion.

4.3 Proper actions

By adding to Section 4.2's main theorem (i.e., R is a closed submanifold of $M \times M$ iff M/G is a manifold with $\pi : M \rightarrow M/G$ a submersion), the notion of a proper action we can give useful sufficient conditions for:

- (A): orbits to be submanifolds;
- (B): M/G to be a manifold.

An action $\Phi : G \times M \rightarrow M$ is called *proper* if the map

$$(147) \quad \tilde{\Phi} : (g, x) \in G \times M \mapsto (x, \Phi(g, x)) \in M \times M$$

is proper. By this we mean that if $\{x_n\}$ is a convergent sequence in M , and $\{\Phi_{g_n}(x_n)\}$ is a convergent sequence in M , then $\{g_n\}$ has a convergent subsequence in G . In finite dimensions, this means that compact sets have compact inverse images; i.e. if $K \subset M \times M$ is compact, then $\tilde{\Phi}^{-1}(K)$ is compact.

If G is compact, this condition is automatically satisfied. Also, the action of a group on itself by left (or by right) translation (Example (iv) of Section 4.1.0.1) is always proper. Furthermore, the cotangent lift of left (or right) translation ((2) and Example (viii) of Section 4.1.0.1) is always proper. We shall not prove this, but it will be important in the sequel.

4.3.0.1 Isotropy groups; orbits as manifolds For $x \in M$ the *isotropy* (or *stabilizer* or *symmetry*) group of Φ at x is

$$(148) \quad G_x := \{g \in G : \Phi_g(x) \equiv \Phi(g, x) = x\} \subset G.$$

(So an action is free iff for all $x \in M$, $G_x = \{e\}$.)

So if we define

$$(149) \quad \Phi^x : G \rightarrow M \quad : \quad \Phi^x(g) := \Phi(g, x)$$

we have: $G_x = (\Phi^x)^{-1}(x)$. (The notation Φ^x is a “cousin” of the notation Φ_g defined in eq. 116.)

So since Φ^x is continuous, G_x is a closed subgroup of G . So, by the result in (2) of Section 3.4.3 (i.e. the result before eq. 97), G_x is a submanifold (as well as Lie subgroup) of G . And if the action is proper, G_x is compact.

Furthermore, the fact that for all $h \in G_x$ we have $\Phi^x(gh) = \Phi_g \circ \Phi_h(x) = \Phi_g(x)$, implies that Φ^x naturally induces a map

$$(150) \quad \tilde{\Phi}^x : [g] = gG_x \in G/G_x \mapsto \Phi_g x \in \text{Orb}(x) \subset M.$$

That is, this map is well-defined. $\tilde{\Phi}^x$ is injective because if $\Phi_g x = \Phi_h x$ then $g^{-1}h \in G_x$, so that $gG_x = hG_x$.

It follows from Section 4.2’s main theorem (i.e., R is a closed submanifold of $M \times M$ iff M/G is a manifold with $\pi : M \rightarrow M/G$ a submersion) that:

- (a) If $\Phi : G \times M \rightarrow M$ is an action and $x \in M$, then $\tilde{\Phi}^x$ defined by eq. 150 is an injective immersion.

Here we recall from Section 3.3.1 that injective immersions need not be embeddings. But:

- (b) If also Φ is proper, the orbit $\text{Orb}(x)$ is a closed submanifold of M and $\tilde{\Phi}^x$ is a diffeomorphism. In other words: the manifold structure of $\text{Orb}(x)$ is given by the bijective map $[g] \in G/G_x \mapsto g \cdot x \in \text{Orb}(x)$ being a diffeomorphism.

Examples:

(We use the numbering of corresponding examples in Section 4.1.0.1):

- (i) $G = SO(3)$ acts on $M = \mathbb{R}^3$ by $(A, x) \mapsto Ax$. Since $\text{Orb}(x)$ is a sphere centred at the origin of radius $\|x\|$, $M/G \cong \mathbb{R}^+$: which is not a manifold. But results (a) and (b) are illustrated: the isotropy group G_x at x is the group of rotations with x on the axis; the action is proper (for G is compact); the orbit $\text{Orb}(x)$ is a closed manifold of M ; and the isotropy group’s cosets $[g] \in G/G_x$ are mapped diffeomorphically by $\tilde{\Phi}^x$ to points on the sphere $\text{Orb}(x)$.

- (iii) Let X be the constant vector field ∂_x on $M = \mathbb{R}^3$. X is complete. The action of \mathbb{R} on M has as orbit through the point $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$, the line $y = \text{constant}$, $z = \text{constant}$. The action is free, and therefore faithful and

the isotropy groups are trivial. So $G/G_x \equiv G$. The action is proper. Again results (a) and (b) are illustrated: the orbits $\text{Orb}(\mathbf{x})$ are closed submanifolds of M , viz. copies of the real line $\mathbb{R} = G \equiv G/G_x$ that are diffeomorphic to \mathbb{R} by $\tilde{\Phi}^{\mathbf{x}}$.

4.3.0.2 A sufficient condition for the orbit space M/G to be a manifold

With result (b) from the end of Section 4.3.0.1, we can prove that:

If $\Phi : G \times M \rightarrow M$ is a proper free action, then the orbit space M/G is a manifold with $\pi : M \rightarrow M/G$ a submersion.

Examples: (again using the numbering in Section 4.1.0.1):

- (i) $G = SO(3)$ acts on $M = \mathbb{R}^3$ by $(A, x) \mapsto Ax$. Since $\text{Orb}(x)$ is a sphere centred at the origin of radius $\|x\|$, $M/G \cong \mathbb{R}^+$: which is not a manifold, and indeed the action is not free.
- (iii) Let X be the constant vector field ∂_x on $M = \mathbb{R}^3$. X is complete, and the action of \mathbb{R} on M has as orbits the lines $y = \text{constant}$, $z = \text{constant}$. The action is faithful, free and proper, so that the orbit space M/G is a manifold: $M/G \cong \mathbb{R}^2$.
- (iv) Left (or right) translation is obviously a free action of a group G on itself, and we noted above that it is proper. But since it is transitive, the orbit space G/G is the trivial 0-dimensional manifold (the singleton set of G).
- (viii) The cotangent lift of left (or right) translation by $SO(3)$, or more generally, by a Lie group G . This action is proper (noted after eq. 147), and obviously free.
- (ix) The Euclidean group E acts freely on $M := T^*\mathbb{R}^{3N} - (\delta \cup \Delta)$. This action is also proper: a (harder!) exercise for the reader.

4.4 Infinitesimal generators of actions

We now connect this Subsection's topic, group actions, with the Lie algebra of the Lie group concerned, i.e. with the topic of Section 3.4, especially 3.4.2.

Let $\Phi : G \times M \rightarrow M$ be a (left) action by the Lie group G on a manifold M . Then each $\xi \in \mathfrak{g}$ defines an action of \mathbb{R} on M , which we write as Φ^ξ , in the following way.

We can think either in terms of exponentiation of ξ 's corresponding left-invariant vector field X_ξ (cf. eq. 36 and 75); or in terms of exponentiating ξ itself (cf. eq. 78 and 79):

$$(151) \quad \Phi^\xi : \mathbb{R} \times M \rightarrow M \quad : \quad \Phi^\xi(\tau, x) := \Phi(\exp(\tau X_\xi), x) \equiv \Phi(\exp(\tau \xi), x).$$

That is, in terms of our subscript notation for the original action Φ (cf. eq. 116): $\Phi_{\exp(\tau X_\xi)} \equiv \Phi_{\exp(\tau\xi)} : M \rightarrow M$ is a flow on M .

That the flow is complete, i.e. that an action of all of \mathbb{R} is defined, follows from (2) *Exponentiation again* of Section 3.4.2, especially after eq. 76. Cf. also example (iii) of Section 4.1.

We say that the corresponding vector field on M , written ξ_M , i.e. the vector field defined at $x \in M$ by

$$(152) \quad \xi_M(x) := \frac{d}{d\tau} \Big|_{\tau=0} \Phi_{\exp(\tau X_\xi)}(x) \equiv \frac{d}{d\tau} \Big|_{\tau=0} \Phi_{\exp(\tau\xi)}(x)$$

is the *infinitesimal generator* of the action corresponding to ξ .

In terms of the map Φ^x defined in eq. 149, we have that for all $\xi \in \mathfrak{g}$

$$(153) \quad \xi_M(x) = (T_e \Phi^x)(\xi) .$$

So NB: the words ‘infinitesimal generator’ are used in different, though related, ways. In Remark (2) at the end of Section 3.4.2, a vector field on the group G , or an element $\xi \in \mathfrak{g}$, was called an ‘infinitesimal generator’. Here the infinitesimal generator is a vector field on the action-space M . Similarly, beware the notation: ξ_M is a vector field on M , while X_ξ is a vector field on G .

As an example, we again take the rotation group $SO(3)$ acting on \mathbb{R}^3 : $(A, \mathbf{x}) \in SO(3) \times \mathbb{R} \mapsto A\mathbf{x}$. One readily checks that with $\omega \in \mathbb{R}^3$, so that $\Theta(\omega) \in \mathfrak{so}(3)$, the infinitesimal generator of the action corresponding to $\xi \equiv \Theta(\omega)$ is the vector field on \mathbb{R}^3

$$(154) \quad \xi_{\mathbb{R}^3}(\mathbf{x}) \equiv (\Theta(\omega))_{\mathbb{R}^3}(\mathbf{x}) = \omega \wedge \mathbf{x} .$$

In particular, the vector field on \mathbb{R}^3 representing infinitesimal anti-clockwise rotation about the x -axis is $e_1 := y\partial_z - z\partial_y$ (cf. eq. 107). Similarly, the infinitesimal generators of the action of rotating about the y axis and about the z -axis are, respectively: $e_2 := z\partial_x - x\partial_z$ and $e_3 := x\partial_y - y\partial_x$. The Lie brackets are given by:

$$(155) \quad [e_1, e_2] = -e_3 \quad [e_3, e_1] = -e_2 \quad [e_2, e_3] = -e_1 .$$

The minus signs here are a general feature of the transition $\xi \in \mathfrak{g} \mapsto \xi_M \in \mathcal{X}(M)$; cf. result (4) below.

As another example, we take the infinitesimal generator of left and right translation on the group G . (We will need this example for our theorems about symplectic reduction; cf. Sections 6.5.3, 7.2 and 7.3.3.) NB: There will be a ‘left-right flip’ here, which continues the discussion in (4) of Section 3.4.2.2, comparing using left-invariant vs. right-invariant vector fields to define the Lie algebra of a Lie group.

For left translation $\Phi(g, h) \equiv L_g h := gh$, we have for all $\xi \in \mathfrak{g}$:

$$(156) \quad \Phi^\xi(\tau, h) = (\exp \tau\xi)h = R_h(\exp \tau\xi) ;$$

so that the infinitesimal generator is

$$(157) \quad \xi_G(g) = (T_e R_g)\xi .$$

So ξ_G is a *right*-invariant vector field; and unless G is abelian, it is *not* equal to the left-invariant vector field $g \mapsto X_\xi(g) := (T_e L_g)\xi$; cf. eq. 68 and 70.

Similarly, for right translation (which is a *right* action, cf. (1) (iv) in Section 4.1.0.1), the infinitesimal generator is the left-invariant vector field

$$(158) \quad g \mapsto X_\xi(g) := (T_e L_g)\xi .$$

Three straightforward results connect the notion of an infinitesimal generator with previous ideas. I will not give proofs, but will present them in the order of the previous ideas.

- (1) Recall the correspondence between Lie subgroups and Lie subalgebras, at the end of Section 3.4.3; eq. 97. This implies that the Lie algebra of the isotropy group $G_x, x \in M$ (called the *isotropy algebra*), is

$$(159) \quad \mathfrak{g}_x = \{\xi \in \mathfrak{g} : \xi_M(x) = 0\} .$$

- (2) Infinitesimal generators ξ_M give a differential version of the notion of equivariance, discussed in (1) of Section 4.2: a version called *infinitesimal equivariance*.

In eq. 144, we set $g = \exp(\tau\xi)$ and differentiate with respect to τ at $\tau = 0$. This gives $Tf \circ \xi_M = \xi_N \circ f$. That is: ξ_M and ξ_N are f -related. In terms of the pullback f^* of f , we have: $f^*\xi_N = \xi_M$.

- (3) Suppose the action Φ is proper, so that by result (b) at the end of Section 4.3.0.1: the orbit $\text{Orb}(x)$ of any point $x \in M$ is a (closed) submanifold of M . Then the tangent space to $\text{Orb}(x)$ at a point y in $\text{Orb}(x)$ is

$$(160) \quad T\text{Orb}(x)_y = \{\xi_M(y) : \xi \in \mathfrak{g}\} .$$

Finally, there is a fourth result relating infinitesimal generators ξ_M to previous ideas; as follows. (But it is less straightforward than the previous (1)–(3): its proof requires the notion of the adjoint representation, described in the next Section.)

- (4) The infinitesimal generator map $\xi \mapsto \xi_M$ establishes a Lie algebra *anti*-homomorphism between \mathfrak{g} and the Lie algebra \mathcal{X}_M of all vector fields on M . (Contrast the Lie algebra isomorphism between \mathfrak{g} and the set $\mathcal{X}_L(G)$ of left-invariant vector fields on the group G ; Section 3.4.2 especially eq. 70.) That is:

$$(161) \quad (a\xi + b\eta)_M = a\xi_M + b\eta_M \quad ; \quad [\xi_M, \eta_M] = -[\xi, \eta]_M \quad \forall \xi, \eta \in \mathfrak{g}, \text{ and } a, b \in \mathbb{R} .$$

Incidentally, returning to (4) of Section 3.4.2.2, which considered defining the Lie algebra of a Lie group in terms of right-invariant vector fields, instead of left-invariant vector fields: had we done so, the corresponding map $\xi \mapsto \xi_M$ would have been a Lie algebra homomorphism.

4.5 The adjoint and co-adjoint representations

A leading idea of later Sections (especially Sections 5.4, 6.4 and 7) will be that there is a natural symplectic structure in the orbits of a certain natural representation of any Lie group: namely a representation of the group on the dual of its own Lie algebra, called the *co-adjoint representation*. Here we introduce this representation. But we lead up to it by first describing the *adjoint representation* of a Lie group on its own Lie algebra. Even apart from symplectic structure (and so applications in mechanics), both representations illustrate the ideas of previous Subsections. I will again use $SO(3)$ and $\mathfrak{so}(3)$ as examples.

4.5.1 The adjoint representation

We proceed in four stages. We first define the representation, then discuss infinitesimal generators, then discuss matrix Lie groups, and finally discuss the rotation group.

(1): *The representation defined:*

Let G be a Lie group and \mathfrak{g} its Lie algebra, i.e. the tangent space to the group at the identity $e \in G$, equipped with the commutator bracket operation $[\cdot, \cdot]$.

Recall (e.g. from the beginning of Section 3.4.2) that G acts on itself by left and right translation: each $g \in G$ defines diffeomorphisms of G onto itself by

$$(162) \quad L_g : h \in G \mapsto gh \in G \quad ; \quad R_g : h \in G \mapsto hg \in G.$$

The induced maps of the tangent spaces are, for each $h \in G$:

$$(163) \quad L_{g*} : TG_h \rightarrow TG_{gh} \quad \text{and} \quad R_{g*} : TG_h \rightarrow TG_{hg}.$$

The diffeomorphism $K_g := R_{g^{-1}} \circ L_g$ (i.e. conjugation by g , $K_g : h \mapsto ghg^{-1}$) is an inner automorphism of G . (Cf. example (v) at the end of Section 4.1.) Its derivative at the identity $e \in G$ is a linear map from the Lie algebra \mathfrak{g} to itself, which is denoted:

$$(164) \quad Ad_g := (R_{g^{-1}} \circ L_g)_{*e} : \mathfrak{g} \rightarrow \mathfrak{g}.$$

So letting g vary through G , the map $Ad : g \mapsto Ad_g$ assigns to each g a member of $\text{End}(\mathfrak{g})$, the space of linear maps on (endomorphisms of) \mathfrak{g} . The chain rule implies that $Ad_{gh} = Ad_g Ad_h$. So

$$(165) \quad Ad : g \mapsto Ad_g$$

is a left action, a representation, of G on \mathfrak{g} : $G \times \mathfrak{g} \rightarrow \mathfrak{g}$. It is called the *adjoint representation*.

Three useful results about Ad follow from our results (1) and (3) in Section 3.4.2.2 (cf. eq. 80: *Homomorphisms respect exponentiation*):

- [1] If $\xi \in \mathfrak{g}$ generates the one-parameter subgroup $H = \{\exp(\tau X_\xi) : \tau \in \mathbb{R}\}$, then $Ad_g(\xi)$ generates the conjugate subgroup $K_g(H) = gHg^{-1}$.

$$(166) \quad \exp(Ad_g(\xi)) = K_g(\exp \xi) := g(\exp \xi)g^{-1}.$$

Incidentally, eq. 166 has a many-parameter generalization. Let H and H' be two connected r -dimensional Lie subgroups of the Lie group G , with corresponding Lie subalgebras \mathfrak{h} and \mathfrak{h}' of the Lie algebra $\mathfrak{g} = \mathfrak{g}(G)$. Then H and H' are conjugate subgroups, $H' = gHg^{-1}$, iff \mathfrak{h} and \mathfrak{h}' are corresponding conjugate subalgebras, i.e. $\mathfrak{h}' = Ad_g(\mathfrak{h})$.

[2] Eq. 166 also implies another result which will be needed for a crucial result about symplectic reduction, in Section 6.5.2. (The many-parameter generalization just mentioned will not be needed.) It relates Ad to the pullback of an arbitrary action Φ .

Thus let Φ be a left action of G on M . Then for every $g \in G$ and $\xi \in \mathfrak{g}$

$$(167) \quad (Ad_g \xi)_M = \Phi_{g^{-1}}^* \xi_M,$$

where Φ^* indicates pullback of the vector field. For we have:

$$(168) \quad (Ad_g \xi)_M(x) := \left. \frac{d}{d\tau} \right|_{\tau=0} \Phi(\exp(\tau Ad_g \xi), x)$$

$$(169) \quad = \left. \frac{d}{d\tau} \right|_{\tau=0} \Phi(g(\exp \tau \xi)g^{-1}, x) \quad \text{by eq. 166}$$

$$(170) \quad = \left. \frac{d}{d\tau} \right|_{\tau=0} (\Phi_g \circ \Phi_{\exp \tau \xi} \circ \Phi_{g^{-1}}(x))$$

$$(171) = T_{\Phi_{g^{-1}}(x)} \Phi_g(\xi_M(\Phi_{g^{-1}}(x))) \quad \text{by the chain rule and eq. 152}$$

$$(172) \quad = \left(\Phi_{g^{-1}}^* \xi_M \right)(x) \quad \text{by the definition of pullback.}$$

Not only is this result needed later. Also, incidentally: it is the main part of the proof of result (4) at the end of Section 4.4, that $\xi \mapsto \xi_M$ is a Lie algebra anti-homomorphism.

[3] Ad_g is an algebra homomorphism, i.e.

$$(173) \quad Ad_g[\xi, \eta] = [Ad_g \xi, Ad_g \eta], \quad \xi, \eta \in \mathfrak{g}.$$

(2): *Infinitesimal generators: the map ad:*

The map Ad is differentiable. Its derivative at $e \in G$ is a linear map from the Lie algebra \mathfrak{g} to the space of linear maps on \mathfrak{g} . This map is called ad , and its value for argument $\xi \in \mathfrak{g}$ is written ad_ξ . That is:

$$(174) \quad ad := Ad_{*e} : \mathfrak{g} \rightarrow \text{End } \mathfrak{g} \quad ; \quad ad_\xi = \left. \frac{d}{d\tau} \right|_{\tau=0} Ad_{\exp(\tau \xi)}$$

where $\exp(\tau \xi)$ is the one-parameter subgroup with tangent vector ξ at the identity. But if we apply the definition eq. 152 of the infinitesimal generator of an action, to the adjoint action Ad , we get that for each $\xi \in \mathfrak{g}$, the generator $\xi_{\mathfrak{g}}$, i.e. a vector field on \mathfrak{g} , is

$$(175) \quad \xi_{\mathfrak{g}} : \eta \in \mathfrak{g} \mapsto \xi_{\mathfrak{g}}(\eta) \in \mathfrak{g} \quad \text{with} \quad \xi_{\mathfrak{g}}(\eta) := \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp(\tau\xi)}(\eta).$$

Comparing eq. 174, we see that ad_{ξ} is just the infinitesimal generator $\xi_{\mathfrak{g}}$ of the adjoint action corresponding to ξ :

$$(176) \quad ad_{\xi} = \xi_{\mathfrak{g}}.$$

We now compute the infinitesimal generators of the adjoint action. It will be crucial to later developments (especially Section 5.4) that these are given by the Lie bracket in \mathfrak{g} .

We begin by considering the function $Ad_{\exp(\tau\xi)}(\eta)$ to be differentiated. By eq. 164, we have

$$(177) \quad \begin{aligned} Ad_{\exp(\tau\xi)}(\eta) &= T_e(R_{\exp(-\tau\xi)} \circ L_{\exp(\tau\xi)})(\eta) \\ &= (T_{\exp(\tau\xi)}(R_{\exp(-\tau\xi)}) \circ T_e L_{\exp(\tau\xi)})(\eta) \\ &= (T_{\exp(\tau\xi)}(R_{\exp(-\tau\xi)}) \cdot X_{\eta}(\exp(\tau\xi))) \end{aligned}$$

where the second line follows by the chain rule, and the third by definition of left-invariant vector field. Writing the flow of X_{ξ} as $\phi_{\tau}(g) = g \exp \tau\xi = R_{\exp(\tau\xi)}g$, and applying the definition of the Lie derivative (eq. 45), we then have

$$(178) \quad \begin{aligned} \xi_{\mathfrak{g}}(\eta) &:= \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp(\tau\xi)}(\eta) = \frac{d}{d\tau} [T_{\phi_{\tau}(e)}\phi_{\tau}^{-1} \cdot X_{\eta}(\phi_{\tau}(e))] \Big|_{\tau=0} \\ &= [X_{\xi}, X_{\eta}](e) = [\xi, \eta]. \end{aligned}$$

where the final equation is the definition eq. 73 of the Lie bracket in the Lie algebra.

So for the adjoint action, the infinitesimal generator corresponding to ξ is taking the Lie bracket: $\eta \mapsto [\xi, \eta]$. To sum up: eq. 174 and 175 now become

$$(179) \quad ad = Ad_{*e} : \mathfrak{g} \rightarrow \text{End } \mathfrak{g} \quad ; \quad ad_{\xi} = \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp(\tau\xi)} = \xi_{\mathfrak{g}} : \eta \in \mathfrak{g} \mapsto [\xi, \eta] \in \mathfrak{g}.$$

(3): *Example: matrix Lie groups:*

In the case where $G \subset GL(n, \mathbb{R})$ is a matrix Lie group with Lie algebra $\mathfrak{g} \subset \mathfrak{gl}(n)$, these results are easy to verify. Writing $n \times n$ matrices as $A, B \in G$, conjugation is $K_A(B) = ABA^{-1}$, and the adjoint map Ad is also given by conjugation

$$(180) \quad Ad_A(X) = AXA^{-1}, \quad A \in G, X \in \mathfrak{g}.$$

So with $A(\tau) = \exp(\tau X)$, so that $A(0) = I$ and $A'(0) = X$, we have with $Y \in \mathfrak{g}$

$$(181) \quad \begin{aligned} \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp \tau X} Y &= \frac{d}{d\tau} \Big|_{\tau=0} [A(\tau)Y A(\tau)^{-1}] \\ &= A'(0)Y A^{-1}(0) + A(0)Y A^{-1}'(0). \end{aligned}$$

But differentiating $A(\tau)A^{-1}(\tau) = I$ yields

$$(182) \quad \frac{d}{d\tau}(A^{-1}(\tau)) = -A^{-1}(\tau)A'(\tau)A^{-1}(\tau), \quad \text{and so } A^{-1'}(0) = -A'(0) = -X$$

so that indeed we have

$$(183) \quad \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp \tau X} Y = XY - YX = [X, Y].$$

(4): *Example: the rotation group:*

It is worth giving details for the case of $G = SO(3)$, $\mathfrak{g} = \mathfrak{so}(3)$. We saw in Section 3.4.4 (eq. 107) that the three matrices

$$(184) \quad A^x = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \quad A^y = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad A^z = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

span $\mathfrak{so}(3)$, and generate the one-parameter subgroups

$$(185) \quad R_\theta^x = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}, \quad R_\theta^y = \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix},$$

$$R_\theta^z = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

representing anticlockwise rotation around the respective coordinate axes in the physical space \mathbb{R}^3 . To calculate the adjoint action of R_θ^x on the generator A^y , we differentiate the product $R_\theta^x R_\tau^y R_{-\theta}^x$ with respect to τ and set $\tau = 0$. That is, we find

$$(186) \quad Ad_{R_\theta^x}(A^y) = R_\theta^x(A^y)R_\theta^x = \begin{pmatrix} 0 & -\sin \theta & \cos \theta \\ \sin \theta & 0 & 0 \\ -\cos \theta & 0 & 0 \end{pmatrix} = \cos \theta \cdot A^y + \sin \theta \cdot A^z.$$

We similarly find

$$(187) \quad Ad_{R_\theta^x}(A^x) = A^x, \quad Ad_{R_\theta^x}(A^z) = -\sin \theta \cdot A^y + \cos \theta \cdot A^z.$$

So the adjoint action of the subgroup R_θ^x representing rotations around the x -axis of physical space is given by rotations around the A^x -axis in the Lie algebra space $\mathfrak{so}(3)$. Similarly for the other subgroups representing rotations around the y or z -axis. And so for any rotation matrix $R \in SO(3)$, relative to given axes x, y, z for \mathbb{R}^3 , its adjoint map Ad_R acting on $\mathfrak{so}(3) \cong \mathbb{R}^3$ has the same matrix representation relative to the induced basis $\{A^x, A^y, A^z\}$ of $\mathfrak{so}(3)$. (NB: This agreement between $SO(3)$'s adjoint representation and its natural physical interpretation is special to $SO(3)$: it does not hold for other matrix Lie groups.)

Finally, the infinitesimal generators of the adjoint action are given by differentiation. For example, using eq. 108, we find that

$$(188) \quad ad_{A^x}(A^y) := \frac{d}{d\theta} \Big|_{\theta=0} Ad_{R_\theta^x} A^y = A^z ;$$

which agrees with the commutator: $A^z = [A^x, A^y]$.

4.5.2 The co-adjoint representation

Again we proceed in stages. We first define the representation, then discuss infinitesimal generators, and then take the rotation group as an example.

(1): *The representation defined:*

We recall that a linear map $A : V \rightarrow W$ induces (basis-independently) a *transpose* (dual) map, written A^* (or \tilde{A} or A^T), $A^* : W^* \rightarrow V^*$ on the dual spaces, $V^* := \{ \alpha : V \rightarrow \mathbb{R} \mid \alpha \text{ linear} \}$ and similarly for W^* ; by

$$(189) \quad \forall \alpha \in W^*, \forall v \in V : \quad A^*(\alpha)(v) \equiv \langle A^*(\alpha) ; v \rangle := \alpha(A(v)) \equiv (\alpha \circ A)(v) .$$

So any representation, \mathcal{R} say, of a group G on a vector space V , $\mathcal{R} : G \rightarrow \text{End}(V)$, induces a representation \mathcal{R}^* of G on the dual space V^* , by taking the transpose. We shall call \mathcal{R}^* the *dual* or *transpose* of \mathcal{R} ; it is also sometimes called a ‘contragredient representation’. That is: for $\mathcal{R}(g) : V \rightarrow V$, we define $\mathcal{R}^*(g) : V^* \rightarrow V^*$ by

$$(190) \quad \mathcal{R}^*(g) : \alpha \in V^* \mapsto \mathcal{R}^*(g)(\alpha) := \alpha(\mathcal{R}(g)) \in V^* .$$

Thus the adjoint representation of G on \mathfrak{g} induces a *co-adjoint representation* of G on the dual \mathfrak{g}^* of its Lie algebra \mathfrak{g} , i.e. on the cotangent space to the group G at the identity, $\mathfrak{g}^* = T_e^*G$. The co-adjoint representation will play a central role in symplectic reduction (starting in Section 5.4).

So let $Ad_g^* : \mathfrak{g}^* \rightarrow \mathfrak{g}^*$ be the dual (aka: transpose) of Ad_g , defined by

$$(191) \quad \forall \alpha \in \mathfrak{g}^*, \xi \in \mathfrak{g} : \quad \langle Ad_g^* \alpha ; \xi \rangle := \langle \alpha ; Ad_g \xi \rangle .$$

Since $Ad : g \mapsto Ad_g$ is a left action ($Ad_{gh} = Ad_g Ad_h$), the assignment $g \mapsto Ad_g^*$ is a right action. So to define a left action, we use the inverse g^{-1} ; cf. eq. 117 and 126. Namely, we define the left action

$$(192) \quad (g, \alpha) \in G \times \mathfrak{g}^* \mapsto Ad_{g^{-1}}^* \alpha \in \mathfrak{g}^* ;$$

called the *co-adjoint action* of G on \mathfrak{g}^* . And the corresponding *co-adjoint representation* of G on \mathfrak{g}^* is denoted by

$$(193) \quad Ad^* : G \rightarrow \text{End}(\mathfrak{g}^*), \quad Ad_{g^{-1}}^* = (T_e(R_g \circ L_{g^{-1}}))^* .$$

(2): *The map ad^* ; infinitesimal generators:*

The map Ad^* is differentiable. Its derivative at $e \in G$ is a linear map from the Lie algebra \mathfrak{g} to the space of linear maps on \mathfrak{g}^* . This map is called ad^* , and its value for argument $\xi \in \mathfrak{g}$ is written ad_ξ^* . Thus ad_ξ^* is an endomorphism of \mathfrak{g}^* , and we have

$$(194) \quad ad^* = Ad_{*e}^* : \xi \in \mathfrak{g} \rightarrow ad_\xi^* \in \text{End} \mathfrak{g}^* .$$

Now recall our deduction from eq. 174 and 175 that $ad_\xi = \xi_{\mathfrak{g}}$, i.e. eq. 176. In the same way we here deduce an equality to the infinitesimal generator of the co-adjoint action:

$$(195) \quad ad_{\xi}^* = \xi_{\mathfrak{g}^*} .$$

In fact, ad_{ξ}^* is, *modulo* a minus sign, the *adjoint* of ad_{ξ} , in the usual sense of the natural pairing of a vector space with its dual: as we now show. (So the notation ad^* is justified, *modulo* a minus sign.)

Let us compute for this action, the value of the infinitesimal generator $\xi_{\mathfrak{g}^*}$ (a vector field on \mathfrak{g}^* , induced by $\xi \in \mathfrak{g}$) at the point $\alpha \in \mathfrak{g}^*$. That is, we will compute the value $\xi_{\mathfrak{g}^*}(\alpha)$. As usual, we identify the tangent space $(T\mathfrak{g}^*)_{\alpha}$ in which this value lives, with \mathfrak{g}^* itself; and similarly for \mathfrak{g} . So, with $\xi_{\mathfrak{g}^*}$ acting on $\eta \in \mathfrak{g}$, we compute:

$$(196) \quad \langle ad_{\xi}^*(\alpha); \eta \rangle \equiv \langle \xi_{\mathfrak{g}^*}(\alpha); \eta \rangle = \left\langle \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp -\tau\xi}^*(\alpha); \eta \right\rangle$$

$$(197) \quad = \frac{d}{d\tau} \Big|_{\tau=0} \langle Ad_{\exp -\tau\xi}^*(\alpha); \eta \rangle = \frac{d}{d\tau} \Big|_{\tau=0} \langle \alpha; Ad_{\exp -\tau\xi} \eta \rangle$$

$$(198) = \left\langle \alpha; \frac{d}{d\tau} \Big|_{\tau=0} Ad_{\exp -\tau\xi} \eta \right\rangle = \langle \alpha; -[\xi, \eta] \rangle = - \langle \alpha; ad_{\xi}(\eta) \rangle .$$

So ad_{ξ}^* , defined as the derivative of Ad^* is, up to a sign, the adjoint of ad_{ξ} .

(3): *Example: the rotation group:*

Let us now write the elementary vector product in \mathbb{R}^3 as \wedge , and identify $\mathfrak{so}(3) \cong (\mathbb{R}^3, \wedge)$ and $\mathfrak{so}(3)^* \cong \mathbb{R}^{3*}$. And let us have the natural pairing given by the elementary euclidean inner product \cdot . Then the result just obtained (now with \bullet marking the argument-place)

$$(199) \quad \langle \xi_{\mathfrak{g}^*}(\alpha); \bullet \rangle = - \langle \alpha; [\xi, \bullet] \rangle$$

becomes for $\alpha \in \mathfrak{so}(3)^*$ and $\xi \in \mathfrak{so}(3)$

$$(200) \quad \xi_{\mathfrak{so}(3)^*}(\alpha) \cdot \bullet = -\alpha \cdot (\xi \wedge \bullet) .$$

So for $\eta \in \mathfrak{so}(3)$, we have

$$(201) \quad \langle \xi_{\mathfrak{so}(3)^*}(\alpha); \eta \rangle = \xi_{\mathfrak{so}(3)^*}(\alpha) \cdot \eta = -\alpha \cdot (\xi \wedge \eta) = -(\alpha \wedge \xi) \cdot \eta = - \langle \alpha \wedge \xi; \eta \rangle .$$

In short:

$$(202) \quad \xi_{\mathfrak{so}(3)^*}(\alpha) = -\alpha \wedge \xi = \xi \wedge \alpha .$$

Now since $SO(3)$ is compact, we know that the co-adjoint action is proper; so $\text{Orb}(\alpha)$ is a closed submanifold of $\mathfrak{so}(3)^*$, and eq. 160 of Section 4.4 applies. So if we fix α , and let ξ vary through $\mathfrak{so}(3) \cong \mathbb{R}^3$, we get all of the tangent space $T_{\alpha}\text{Orb}(\alpha)$ to the orbit passing through α . Then eq. 202 implies that the tangent space is the plane normal to α , and passing through α 's end-point. Letting α vary through $\mathfrak{so}(3)^*$, we conclude that the *co-adjoint orbits are the spheres centred on the origin*.

In the following Sections, we will see that the orbits of the co-adjoint representation of *any* Lie group G have a natural symplectic structure. So the orbits are

always even-dimensional; and by considering all Lie groups and all possible orbits, we can get a series of examples of symplectic manifolds.

Besides, this fact will play a central role in our generalized formulation of Hamiltonian mechanics, and in symplectic reduction. And we will (mercifully!) get a good understanding of that role, already in Section 5.1. To prepare for that, it is worth gathering some threads about our recurrent example, $SO(3)$; and generalizing them to other Lie groups ...

4.6 Kinematics on Lie groups

To summarize some aspects of this Section, and to make our later discussion of reduction clearer, it is worth collecting and generalizing some of our results about $SO(3)$ and the description it provides of the rigid body. More precisely, we will now combine:

- (i) the description of space and body coordinates in terms of left and right translation, at the end of Section 3.4.4;
- (ii) the cotangent lift of translation (example (viii) of Section 4.1.0.1);
- (iii) the adjoint and co-adjoint representations of $SO(3)$ (as in (4) of Section 4.5.1, and (3) of Section 4.5.2).

We will also generalize: namely, we will consider (i) to (iii) for an arbitrary Lie group G , not just for $SO(3)$. (The point of doing so will become clear in (3) of Section 5.1.) This will occur already in Section 4.6.1. Then in Section 4.6.2, we will show how this material yields natural diffeomorphisms

$$(203) \quad TG \rightarrow G \times \mathfrak{g} \quad \text{and} \quad T^*G \rightarrow G \times \mathfrak{g}^* \quad ;$$

(so if $\dim G = n$, then all four manifolds are $2n$ -dimensional). We will also see that by applying Section 4.2's notion of equivariance, we can "pass to the quotients", and get from eq. 203, the natural diffeomorphisms

$$(204) \quad TG/G \rightarrow \mathfrak{g} \quad \text{and} \quad T^*G/G \rightarrow \mathfrak{g}^* \quad ;$$

where the quotients on the left hand sides (the domains) is by the action of left translation; (to be precise: by the action of its derivative for TG , and its cotangent lift for T^*G).

4.6.1 Space and body coordinates generalized to G

So let a (finite-dimensional) Lie group G act on itself by left and right translation, L_g and R_g . For any $g \in G$, we define

$$(205) \quad \lambda_g : T_g G \rightarrow \mathfrak{g} \quad \text{by} \quad v \in T_g G \mapsto (T_e L_g)^{-1}(v) \equiv (T_g L_{g^{-1}})(v) \in \mathfrak{g} .$$

We similarly define

$$(206) \quad \rho_g : v \in T_g G \mapsto (T_e R_g)^{-1}(v) \equiv (T_g R_{g^{-1}})(v) \in \mathfrak{g} .$$

On analogy with the case of the pivoted rigid body (cf. eq. 113 and 114, or eq. 137), we say that λ_g represents $v \in T_g G$ in *body coordinates*, and ρ_g represents v in *space coordinates*. We also speak of body and space *representations*. The transition from body to space coordinates is then an isomorphism of \mathfrak{g} ; viz. by eq. 164

$$(207) \quad \forall \xi \in \mathfrak{g}, \quad (\rho_g \circ \lambda_g^{-1})(\xi) = \rho_g(T_e L_g(\xi)) \equiv Ad_g \xi .$$

So we can combine the S and B superscript notation of eq. 137 with Section 4.5.1's notion of the adjoint representation, and write

$$(208) \quad v^S = Ad_g v^B .$$

In a similar way, the cotangent lifts of left and right translation provide isomorphisms between the dual spaces $T_g^* G, g \in G$ and \mathfrak{g}^* . Thus for any $g \in G$, we define

$$(209) \quad \bar{\lambda}_g : T_g^* G \rightarrow \mathfrak{g}^* \quad \text{by} \quad \alpha \in T_g^* G \mapsto \alpha \circ T_e L_g \equiv (T_e L_g)^*(\alpha) \equiv (T_e^* L_g)(\alpha) \in \mathfrak{g}^* ;$$

and similarly

$$(210) \quad \bar{\rho}_g : \alpha \in T_g^* G \mapsto \alpha \circ T_e R_g \equiv (T_e^* R_g)(\alpha) \in \mathfrak{g}^* .$$

And we again use the S and B superscript notation of eq. 137, and define for $\alpha \in T_g^* G$

$$(211) \quad \alpha^S := (T_e^* R_g)(\alpha) \equiv \bar{\rho}_g(\alpha) \quad \text{and} \quad \alpha^B := (T_e^* L_g)(\alpha) \equiv \bar{\lambda}_g(\alpha) ,$$

which are called the *space* (or 'spatial') and *body* representations, respectively, of α . The transition from body to space representations is now an isomorphism of \mathfrak{g}^* ; viz.

$$(212) \quad \forall \alpha \in \mathfrak{g}^*, \quad (\bar{\rho}_g \circ \bar{\lambda}_g^{-1})(\alpha) = Ad_{g^{-1}}^*(\alpha) , \quad \text{i.e.} \quad \alpha^S = Ad_{g^{-1}}^*(\alpha^B) .$$

4.6.2 Passage to the quotients

For later purposes, we need to develop the details of how the element $g \in G$ "carries along throughout" in eq. 205 to 212. More precisely, we have two isomorphisms:

$$(213) \quad TG \cong G \times \mathfrak{g} \quad \text{and} \quad T^*G \cong G \times \mathfrak{g}^* .$$

These are isomorphisms of vector bundles; but we shall not develop the language of fibre bundles. What matters for us is that once we exhibit these isomorphisms, we will see that we have equivariant maps relating two group actions, in the sense of eq. 144 and 145. And this will mean that we can pass to the quotients to infer that TG/G is diffeomorphic to \mathfrak{g} , and correspondingly that T^*G/G is diffeomorphic to \mathfrak{g}^* .

This last diffeomorphism will form the first part of Section 7's main theorem, the Lie-Poisson reduction theorem, which says that T^*G/G and \mathfrak{g}^* are isomorphic as Poisson manifolds. In Section 5 onwards, we will develop the notion of a Poisson manifold, and the significance of this isomorphism for the reduction of mechanical problems.

I should note here that there is a parallel story about the first diffeomorphism, i.e. about TG/G being diffeomorphic to \mathfrak{g} . It forms the first part of another reduction theorem, which is the Lagrangian analogue of Section 7's Lie-Poisson theorem. But since this Chapter has adopted the Hamiltonian approach, I will not go into details. They can be found in Marsden and Ratiu [1999, Sections 1.2, 13.5, 13.6], under the title 'Euler-Poincaré reduction'.

Thus corresponding to eq. 205, we define the isomorphism

$$(214) \quad \lambda : TG \rightarrow G \times \mathfrak{g} \quad \text{by} \quad \lambda(v) := (g, (T_e L_g)^{-1}(v)) \equiv (g, (T_g L_{g^{-1}})(v))$$

with $v \in T_g G$, i.e. $g = \pi_G(v)$ and $\pi_G : TG \rightarrow G$ the canonical projection. (As mentioned concerning eq. 121, it is harmless to (follow many presentations and) conflate a point in TG , i.e. strictly speaking a pair (g, v) , $g \in G$, $v \in T_g G$, with its vector v .) And corresponding to eq. 206, we define the isomorphism

$$(215) \quad \rho : TG \rightarrow G \times \mathfrak{g} \quad \text{by} \quad \rho(v) := (g, (T_e R_g)^{-1}(v)) \equiv (g, (T_g R_{g^{-1}})(v)) .$$

The transition from body to space representations given by eq. 207 now implies

$$(216) \quad (\rho \circ \lambda^{-1})(g, \xi) = \rho(g, T_e L_g(\xi)) = (g, (T_e R_g)^{-1} \circ T_e L_g(\xi)) = (g, Ad_g \xi) .$$

In a similar way, the cotangent bundle T^*G is isomorphic in two ways to $G \times \mathfrak{g}^*$: namely by

$$(217) \quad \bar{\lambda}(\alpha) := (g, \alpha \circ T_e L_g) \equiv (g, (T_e^* L_g)\alpha) \in G \times \mathfrak{g}^* ,$$

and by

$$(218) \quad \bar{\rho}(\alpha) := (g, \alpha \circ T_e R_g) \equiv (g, (T_e^* R_g)\alpha) \in G \times \mathfrak{g}^*$$

where $\alpha \in T_g^* G$, i.e. $g = \pi_G^*(\alpha)$ with $\pi_G^* : T^*G \rightarrow G$ the canonical projection. (Again, we harmlessly conflate a point (g, α) in T^*G with its form $\alpha \in T_g^* G$.)

Let us now compute in the body representation, the actions of: (i) the (derivative of the) left translation map, TL_g , and (ii) the corresponding cotangent lift T^*L_g . This will show that λ and $\bar{\lambda}$ are equivariant maps for certain group actions.

(i) We compute:

$$(219) \quad \begin{aligned} (\lambda \circ TL_g \circ \lambda^{-1})(h, \xi) &= (\lambda \circ TL_g)(h, TL_h(\xi)) = \\ &= \lambda(gh, (TL_g \circ TL_h)(\xi)) = \\ &= (gh, ((TL_{gh})^{-1} \circ TL_{gh})(\xi)) = (gh, \xi) . \end{aligned}$$

So in the body representation, left translation does not act on the vector component. (That is intuitive, in that the vector ξ is "attached to the

body” and so should not vary relative to coordinates fixed in it.) Eq. 219 means that λ is an equivariant map relating left translation TL_g on TG to the G -action on $G \times \mathfrak{g}$ given just by left translation on the first component:

$$(220) \quad \Phi_g((h, \xi)) \equiv g \cdot (h, \xi) := (gh, \xi) .$$

Equivariance means that λ induces a map $\hat{\lambda}$ on the quotients. That is: as in eq. 146, the map

$$(221) \quad \hat{\lambda} : TG/G \rightarrow (G \times \mathfrak{g})/G$$

defined as mapping, for any g , the orbit of any $v \in T_g G$ to the orbit of $\lambda(v)$, i.e.

$$(222) \quad \begin{aligned} \hat{\lambda} : \text{Orb}(v) &\equiv \{u \in TG \mid T_g L_h(v) = u, \text{ some } h \in G\} \mapsto \text{Orb}(\lambda(v)) \\ &\equiv \{(hg, (T_e L_g)^{-1}(v)) \mid \text{some } h \in G\} \end{aligned}$$

is well-defined, i.e. independent of the chosen representative v of the orbit.

Besides, since the canonical projections, $v \in TG \mapsto \text{Orb}(v) \in TG/G$ and $(g, \xi) \mapsto \text{Orb}((g, \xi)) \in (G \times \mathfrak{g})/G$, are submersions, we can apply result (1) of Section 4.2 and conclude that $\hat{\lambda}$ is smooth.

Finally, we notice that since the action of left translation is transitive, we can identify each orbit of the Φ of eq. 220 with its right component $\xi \in \mathfrak{g}$; and so we can identify the set of orbits $(G \times \mathfrak{g})/G$ with \mathfrak{g} .

To sum up: we have shown that TG/G and $(G \times \mathfrak{g})/G$, i.e. in effect \mathfrak{g} , are diffeomorphic:

$$(223) \quad \hat{\lambda} : TG/G \rightarrow (G \times \mathfrak{g})/G \equiv \mathfrak{g} .$$

- (ii) The results for the cotangent bundle are similar to those in (i). On analogy with eq. 219, the action of the cotangent lift of left translation T^*L_g is given in body representation by applying eq. 217 to get

$$(224) \quad (\bar{\lambda} \circ (T^*L_g) \circ \bar{\lambda}^{-1})(h, \alpha) = (g^{-1}h, \alpha) ;$$

or equivalently, now taking the cotangent lift of left translation to define a left action (cf. eq. 126),

$$(225) \quad (\bar{\lambda} \circ (T^*L_{g^{-1}}) \circ \bar{\lambda}^{-1})(h, \alpha) = (gh, \alpha) .$$

So in body representation, left translation does not act on the covector component; (again, an intuitive result in so far as α is “attached to the body”). So eq. 225 means that $\bar{\lambda}$ is an equivariant map relating the cotangent lifted left action of left translation on T^*G to the G -action on $G \times \mathfrak{g}^*$ given just by left translation on the first component:

$$(226) \quad \Phi_g((h, \alpha)) \equiv g \cdot (h, \alpha) := (gh, \alpha) .$$

So, on analogy with eq. 221 and 222, we can pass to the quotients, defining a map

$$(227) \quad \hat{\lambda} : T^*G/G \rightarrow (G \times \mathfrak{g}^*)/G$$

by requiring that for $\alpha \in T_g^*G$, so that $T^*L_{h^{-1}}\alpha \in T_{hg}^*G$:

$$(228) \quad \begin{aligned} \hat{\lambda} : \text{Orb}(\alpha) &\equiv \{\beta \in T^*G \mid \beta = T^*L_{h^{-1}}(\alpha), \\ \text{some } h \in G\} &\mapsto \text{Orb}(\bar{\lambda}(\alpha)) \equiv \{(hg, (T_e^*L_g)(\alpha)) \mid \text{some } h \in G\} \equiv \\ &\quad \{(h, (T_e^*L_g)\alpha) \mid \text{some } h \in G\}. \end{aligned}$$

And finally, we identify the set of orbits $(G \times \mathfrak{g}^*)/G$ with \mathfrak{g}^* , so that we conclude that T^*G/G and \mathfrak{g}^* are diffeomorphic. That is, we think of the diffeomorphism $\hat{\lambda}$ as mapping T^*G/G to \mathfrak{g}^* :

$$(229) \quad \hat{\lambda} : \text{Orb}(\alpha) \equiv \{\beta \in T^*G \mid \beta = T^*L_{h^{-1}}(\alpha), \text{ some } h \in G\} \in T^*G/G \mapsto (T_e^*L_g)(\alpha) \in \mathfrak{g}^*.$$

As I said above, this diffeomorphism is the crucial first part of Section 7's main reduction theorem. But we will see its role there, already in (3) of Section 5.1.

Finally, a result which will *not* be needed later. To calculate the derivatives and cotangent lifts of left translation in *space* representation, we replace λ and $\bar{\lambda}$ by ρ and $\bar{\rho}$ as defined by eq. 215 and 218. We get as the analogues of eq. 219 and 224 respectively:

$$(230) \quad (\rho \circ TL_g \circ \rho^{-1})(h, \xi) = (gh, Ad_g(\xi)),$$

and

$$(231) \quad (\bar{\rho} \circ T^*L_g \circ \bar{\rho}^{-1})(h, \alpha) = (g^{-1}h, Ad_g^*(\alpha)).$$

Though these results are not needed later, they are also analogues of some later results, eq. 399 and 400, which we will need. (Note that, in accordance with the discussion between eq. 191 and 192, eq. 231 involves right actions.)

5 POISSON MANIFOLDS

5.1 Preamble: three reasons for Poisson manifolds

Now that we are equipped with Sections 3 and 4's toolbox of modern geometry, we can develop, in this Section and the two to follow, the theory of symplectic reduction. This Section develops the general theory of Poisson manifolds, as a framework for a generalized Hamiltonian mechanics. Its main results concern the foliation, and quotienting, of Poisson manifolds. Then Section 6 returns us to symmetries and conserved quantities: topics which are familiar from Section 2.1.3, but which Section 6 will discuss in the generalized framework using the notion of a momentum map. Finally, in Section 7 all the pieces of our jigsaw puzzle will come together, in our symplectic reduction theorem.

We already glimpsed in (1) of Section 2.2 the idea of a Poisson manifold as a generalization of a symplectic manifold, that provides the appropriate framework for a generalized Hamiltonian mechanics. It is a manifold equipped with a bracket, called a ‘Poisson bracket’, that has essentially the same formal defining properties as in symplectic geometry except that it can be “degenerate”. In particular, the dimension m of a Poisson manifold M can be even or odd. As we will see, Hamiltonian mechanics can be set up on Poisson manifolds, in a natural generalization of the usual formalism: there are m first-order ordinary differential equations for the time evolution of local coordinates x^1, \dots, x^m , and the time-derivative of any dynamical variable (scalar function on the Poisson manifold M) is given by its Poisson bracket with the Hamiltonian. Besides, this generalization reduces to the usual formalism in the following sense. Any Poisson manifold M is foliated into symplectic manifolds, and any Hamiltonian mechanics of our generalized kind defined on M restricts on each symplectic leaf to a conventional Hamiltonian mechanics using the induced symplectic form.

This last point, the invariance of the symplectic leaves under the dynamics, prompts the question ‘why bother with the Poisson manifold, since the dynamics can be written down on each leaf?’. There are three reasons. I will just mention the first; the rest of Section 5 will develop the second; and the two subsequent Sections will develop the third.

(1): *Parameters and stability:*

The first two reasons concern the fact that for many problems in Hamiltonian mechanics, it is natural to consider an odd-dimensional state-space. One principal way this happens is if the system is characterized by some odd number, say s (maybe $s = 1$), of parameters that are constant in time. Then even though for a fixed value of the parameter(s), there is a Hamiltonian mechanics on a symplectic manifold, of dimension $2n$ say, it is useful to envisage the $2n + s$ dimensional space in order to keep track of how the behaviour of systems depends on the parameters. For example, this is very useful for analysing stability, especially if one can somehow control the value of the parameters. Stability theory (and related fields such as bifurcation theory) are crucially important, and vast, topics — which I will not go into.²⁰

(2): *Odd-dimensional spaces: the rigid body again:*

Secondly, even in the absence of such controllable parameters, there are mechanical systems whose description leads naturally to an odd-dimensional state-space. The paradigm elementary example is the rigid body pivoted at a point (mentioned in (3) of Section 2.2). An elementary analysis, repeated in every textbook, leads to a description of the body by the three components of the angular momentum (relative to body coordinates, i.e. coordinates fixed in the body): these compo-

²⁰Except to note a broad philosophical point. These parameters illustrate the modal or counterfactual involvements of mechanics. The s dimensions of the state-space, and the mathematical constructions built on them, show how rich and structured these involvements are. For a detailed discussion of the modal involvements of mechanics, cf. [Butterfield, 2004].

nents evolve according to the three first-order Euler equations.

This situation prompts two foundational questions; (which of course most textbooks ignore!). First, we note that a configuration of the body is given by three real numbers: viz. to specify the rotation required to rotate the body into the given configuration, from a fiducial configuration. So a conventional Hamiltonian description of the rigid body would use six first-order equations. (Indeed, similarly for a Lagrangian description, if we treat the three q s as variables.) So how is the description by Euler's equations related to a six-dimensional Hamiltonian (or indeed Lagrangian) description?

Second, can the description by the Euler equations be somehow regarded as itself Hamiltonian, or Lagrangian?

This Chapter will not pursue these questions about the rigid body; for details, cf. the references at the end of (3) of Section 2.2. For us, the important point is that the theory of symplectic reduction shows that the answer to the second question is Yes. Indeed, a “resounding Yes”. For we will see very soon (in Section 5.2.4.1) that the three-dimensional space of the components, in body coordinates, of the angular momentum is our prototype example of a Poisson manifold; and the evolution by Euler's equations is the Hamiltonian mechanics on each symplectic leaf of this manifold. In short: in our generalized framework, Euler's equations are *already* in Hamiltonian form.

Furthermore, this Poisson manifold is already familiar: it is $\mathfrak{so}(3)^*$, the dual of the Lie algebra of the rotation group. Here we connect with several previous discussions (and this Chapter's second motto).

First: we connect with the discussion of rotation in Relationist and Reductionist mechanics (Sections 2.3.3 to 2.3.5). In particular, cf. comment (iii) about γ , the three variables encoding the total angular momentum of the system, at the end of Section 2.3.4. (So as regards (1)'s idea of labelling the symplectic leaves by parameters constant in time: in this example, it is the magnitude L of the total body angular momentum which is the parameter.)

Second: we connect with Section 3.4.4's discussion of $\mathfrak{so}(3)$, with Section 4.5.2's discussion of the co-adjoint representation on $\mathfrak{so}(3)^*$, and with Section 4.6's discussion of kinematics on an arbitrary Lie group. As regards the rigid body, the main physical idea is that the action of $SO(3)$ on itself by left translation is interpreted in terms of the coordinate transformation, i.e. rotation, between the space and body coordinate systems.

But setting aside the rigid body: recall that in Section 4.5.2 we saw that for $\mathfrak{so}(3)^*$, the co-adjoint orbits are the spheres centred on the origin. I also announced that they have a natural symplectic structure — and that this was true for the orbits of the co-adjoint representation of *any* Lie group. Now that we have the notion of a Poisson manifold, we can say a bit more, though of course the proofs are yet to come: —

For any Lie group G , the dual of its Lie algebra \mathfrak{g}^* is a Poisson manifold; and G has on \mathfrak{g}^* a co-adjoint representation, whose orbits are the symplectic leaves of \mathfrak{g}^* as a Poisson manifold.

In particular, we remark that the theory of the rigid body just sketched is independent of the dimension of physical space being three: it carries over to $\mathfrak{so}(n)^*$ for any n . So we can readily do the Hamiltonian mechanics of the rigid body in arbitrary dimensions. That sounds somewhat academic! But it leads to a more general point, which is obviously of vast practical importance.

In engineering we often need to analyse or design bodies consisting of two or more rigid bodies jointed together, e.g. at a universal joint. Often the configuration space of such a jointed body can be given by a sequence of rotations (in particular about the joints) and-or translations from a fiducial configuration; so that we can take an appropriate Lie group G as the body's configuration space. If so, we can try to mimic our strategy for the rigid body, i.e. to apply the result just announced. And indeed, for such bodies, the action of left translation, and so the adjoint and co-adjoint representations of G on \mathfrak{g} and \mathfrak{g}^* , can often be physically significant.

But leaving engineering aside, let us sum up this second reason for Poisson manifolds as follows. For some mechanical systems the natural state-space for a Hamiltonian mechanics is a Poisson manifold. And in the paradigm case of the rigid body, there is a striking interpretation of the Poisson manifold's leaves as the orbits of the co-adjoint representation of the rotation group $SO(3)$.

(3): *Reduction:*

My first two reasons have not mentioned reduction. But unsurprisingly, they have several connections with the notion. Here I shall state just one main connection, which links Section 4.6's kinematics on Lie groups to our main reduction theorem: this will be my third motivation for studying Poisson manifolds.

In short, the connection is that:

- (i) For various systems, the configuration space is naturally taken to be a Lie group G ; (as we have just illustrated with the rigid body).
- (ii) So it is natural to set up an orthodox Hamiltonian mechanics of the system on the cotangent bundle T^*G . But (as in the Reductionist procedure of Section 2.3.4) it is also natural to quotient by the lift to the cotangent bundle of G 's action on itself by left translation.
- (iii) When we do this, the resulting reduced phase space T^*G/G is a Poisson manifold. Indeed it is an isomorphic copy of \mathfrak{g}^* . That is, we have an isomorphism of Poisson manifolds: $\mathfrak{g}^* \cong T^*G/G$. This is the Lie-Poisson reduction theorem.

I shall give a bit more detail about each of (i)–(iii).

- (i) For various systems, any configuration can be obtained by acting with an element of the Lie group G on some reference configuration which can itself be labelled by an element of G , say the identity $e \in G$. So we take the Lie group G to be the configuration space. As mentioned in (3) of Section 2.2, there is even an infinite-dimensional example of this: the ideal fluid.

- (ii) So T^*G is the conventional Hamiltonian phase space of the system. But G acts on itself by left translation. We can then consider the quotient of T^*G by the cotangent lift of left translation. Intuitively, this is a matter of “rubbing out” the way that T^*G encodes (i)’s choice of reference configuration. By passing to the quotients as in Section 4.6, we infer that T^*G/G is a manifold. But of course it is in general *not* even-dimensional. For its dimension is $\frac{1}{2}\dim(T^*G) \equiv \dim(G)$. So consider any odd-dimensional G : for example, our old friend, the three-dimensional rotation group $SO(3)$.
- (iii) But T^*G/G is always a Poisson manifold. And it is always isomorphic as a Poisson manifold to \mathfrak{g}^* , with its symplectic leaves being the co-adjoint orbits of \mathfrak{g}^* : $\mathfrak{g}^* \cong T^*G/G$.

I end this third reason for studying Poisson manifolds with two remarks about examples.

The first remark echoes the end of Section 4.5.2, where I said that by considering all possible Lie groups and all the orbits of their co-adjoint representations, we get a series of examples of symplectic manifolds. We can now put this together with the notion of a Poisson manifold, and with the comment at the end of Section 3.4.3, that every (finite-dimensional) Lie algebra is the Lie algebra of a Lie group. In short: we get a series of examples of Poisson manifolds, in either of two equivalent ways: from the dual \mathfrak{g}^* of any (finite-dimensional) Lie algebra \mathfrak{g} ; or from the quotient T^*G/G of the cotangent lift of left translation. In either case, the example is the co-adjoint representation.

The second remark is that there are yet other examples of Poisson manifolds and reductions. Indeed, we noted one in Section 2.3.4: viz. the Reductionist’s reduced phase space $\bar{M} := M/E$, obtained by quotienting the phase space $M := T^*\mathbb{R}^{3N} - (\delta \cup \Delta)$ by the (cotangent lift) of the action of the euclidean group E on \mathbb{R}^{3N} . But I shall not go into further details about this example; (for which cf. the Belot papers listed in Section 2.3.1, and references therein). Here it suffices to note that this example is not of the above form: \mathbb{R}^{3N} is not E , and the action of E on \mathbb{R}^{3N} is not left translation. This of course echoes my remarks at the end of Section 1.2 that the theory of symplectic reduction is too large and intricate for this Chapter to be more than an “appetizer”.

So much by way of motivating Poisson manifolds. The rest of this Section will cover reasons (1) and (2); but reason (3), about reduction, is postponed to Sections 6 and 7. We give some basics about Poisson manifolds, largely in coordinate-dependent language, in Section 5.2. In Section 5.3, we move to a more coordinate-independent language and show that Poisson manifolds are foliated into symplectic manifolds. In Section 5.4, we show that the leaves of the foliation of a finite-dimensional Lie algebra \mathfrak{g}^* are the orbits of the co-adjoint representation of G on \mathfrak{g}^* . Finally in Section 5.5, we prove a general theorem about quotienting a Poisson manifold by the action of Lie group, which will be important for Section 7’s main theorem.

5.2 Basics

In Sections 5.2.1 to 5.2.3, we develop some basic definitions and results about Poisson manifolds. This leads up to Section 5.2.4, where we see that the dual of any finite-dimensional Lie algebra has a natural (i.e. basis-independent) Poisson manifold structure. Throughout, there will be some obvious echoes of previous discussions of anti-symmetric forms, Poisson brackets, Hamiltonian vector fields and Lie brackets (Sections 2.1 and 3.2). But I will for the most part *not* articulate these echoes.

5.2.1 Poisson brackets

A manifold M is called a *Poisson manifold* if it is equipped with a *Poisson bracket* (also known as: *Poisson structure*). A *Poisson bracket* is an assignment to each pair of smooth real-valued functions $F, H : M \rightarrow \mathbb{R}$, of another such function, denoted by $\{F, H\}$, subject to the following four conditions:

(a) Bilinearity:

$$(232) \quad \{aF+bG, H\} = a\{F, H\}+b\{G, H\}; \quad \{F, aG+bH\} = a\{F, G\}+b\{F, H\} \quad \forall a, b \in \mathbb{R}.$$

(b) Anti-symmetry:

$$(233) \quad \{F, H\} = -\{H, F\}.$$

(c) Jacobi identity:

$$(234) \quad \{\{F, H\}, G\} + \{\{G, F\}, H\} + \{\{H, G\}, F\} = 0.$$

(d) Leibniz' rule:

$$(235) \quad \{F, H \cdot G\} = \{F, H\} \cdot G + H \cdot \{F, G\}.$$

In other words: M is a Poisson manifold iff both: (i) the set $\mathcal{F}(M)$ of smooth scalar functions on M , equipped with the bracket $\{, \}$, is a Lie algebra; and (ii) the bracket $\{, \}$ is a derivation in each factor.

Any symplectic manifold is a Poisson manifold. The Poisson bracket is defined by the manifold's symplectic form; cf. eq. 18.

“Canonical” Example:

Let $M = \mathbb{R}^m, m = 2n + l$, with standard coordinates $(q, p, z) = (q^1, \dots, q^n, p^1, \dots, p^n, z^1, \dots, z^l)$. Define the Poisson bracket of any two functions $F(q, p, z), H(q, p, z)$ by

$$(236) \quad \{F, H\} := \sum_i^n \left(\frac{\partial F}{\partial q^i} \frac{\partial H}{\partial p^i} - \frac{\partial F}{\partial p^i} \frac{\partial H}{\partial q^i} \right).$$

Thus this bracket ignores the z coordinates; and if l were equal to zero, it would be the standard Poisson bracket for \mathbb{R}^{2n} as a symplectic manifold. We can immediately deduce the Poisson brackets for the coordinate functions. Those for the q s and p s are as for the usual symplectic case:

$$(237) \quad \{q^i, q^j\} = 0 \quad \{p^i, p^j\} = 0 \quad \{q^i, p^j\} = \delta_{ij}.$$

On the other hand, all those involving the z s vanish:

$$(238) \quad \{q^i, z^j\} = \{p^i, z^j\} = \{z^i, z^j\} \equiv 0.$$

Besides, any function F depending only on the z 's, $F \equiv F(z)$ will have vanishing Poisson brackets with all functions $H : \{F, H\} = 0$.

This example seems special in that M is foliated into $2n$ -dimensional symplectic manifolds, each labelled by l constant values of the z s. But Section 5.3.4 will give a generalization for Poisson manifolds of Darboux's theorem (mentioned at the end of Section 2.1.1): a generalization saying, roughly speaking, that every Poisson manifold "looks locally like this".

For any Poisson manifold, we say that a function $F : M \rightarrow \mathbb{R}$ is *distinguished* or *Casimir* if its Poisson bracket with all smooth functions $H : M \rightarrow \mathbb{R}$ vanishes identically: $\{F, H\} = 0$.

5.2.2 Hamiltonian vector fields

Given a smooth function $H : M \rightarrow \mathbb{R}$, consider the map on smooth functions: $F \mapsto \{F, H\}$. The fact that the Poisson bracket is bilinear and obeys Leibniz's rule implies that this map $F \mapsto \{F, H\}$ is a derivation on the space of smooth functions, and so determines a vector field on M ; (cf. (ii) of Section 3.1.2.2). We call this vector field the *Hamiltonian vector field* associated with (also known as: generated by) H , and denote it by X_H .

But independently of the Poisson structure, the action of any vector field X_H on a smooth function F , $X_H(F)$, also equals $L_{X_H}(F) \equiv dF(X_H)$; (cf. eq. 40). So we have for all smooth F

$$(239) \quad L_{X_H}(F) \equiv dF(X_H) \equiv X_H(F) = \{F, H\}.$$

The equations describing the flow of X_H are called *Hamilton's equations*, for the choice of H as "Hamiltonian".

In the previous example with $M = \mathbb{R}^{2n+l}$, we have

$$(240) \quad X_H = \sum_i^n \left(\frac{\partial H}{\partial p^i} \frac{\partial}{\partial q^i} - \frac{\partial H}{\partial q^i} \frac{\partial}{\partial p^i} \right),$$

and the flow is given by the ordinary differential equations

$$(241) \quad \frac{dq^i}{dt} = \frac{\partial H}{\partial p^i} \quad \frac{dp^i}{dt} = -\frac{\partial H}{\partial q^i} \quad \frac{dz^j}{dt} = 0. \quad i = 1, \dots, n; \quad j = 1, \dots, l.$$

Again, the z s, and any function $F(z)$ solely of them, are distinguished and have a vanishing Hamiltonian vector field. On the other hand, the coordinate functions q^i and p^i generate the Hamiltonian vector fields $-\frac{\partial}{\partial p^i}$ and $\frac{\partial}{\partial q^i}$ respectively.

Two further remarks about eq. 239:

- (1) It follows that a function H is distinguished (i.e. has vanishing Poisson brackets with all functions) iff its Hamiltonian vector field X_H vanishes everywhere. And since the Poisson bracket is antisymmetric, this is so iff H is constant along the flow of all Hamiltonian vector fields.
- (2) This equation is the beginning of the theory of constants of the motion (first integrals), and of Noether's theorem, for Poisson manifolds; just as the corresponding equation was the beginning for the symplectic case. This will be developed in Section 6.

Poisson brackets and Lie brackets:

With the definition eq. 239 in hand, we can readily establish our first important connection between Poisson manifolds and Section 3's Lie structures. Namely: result (2) at the end of Section 3.2.2, eq. 60, is also valid for Poisson manifolds.

That is: the Hamiltonian vector field of the Poisson bracket of scalars F, H on a Poisson manifold M is, upto a sign, the Lie bracket of the Hamiltonian vector fields, X_F and X_H , of F and H :

$$(242) \quad X_{\{F, H\}} = -[X_F, X_H] = [X_H, X_F] .$$

The proof is exactly as for eq. 60.

So the Hamiltonian vector fields, with the Poisson bracket, form a Lie subalgebra of the Lie algebra \mathcal{X}_M of all vector fields on the Poisson manifold M . This result will be important in Section 5.3.3's proof that every Poisson manifold is a disjoint union of symplectic manifolds.

5.2.3 Structure functions

We show that to compute the Poisson bracket of any two functions given in some local coordinates $\mathbf{x} = x^1, \dots, x^m$, it suffices to know the Poisson brackets of the coordinates. For any function $H : M \rightarrow \mathbb{R}$, let the components of its Hamiltonian vector field in the coordinate system \mathbf{x} be written as $h^i(x)$. So $X_H = \sum_i^m h^i(x) \frac{\partial}{\partial x^i}$. Then for any other function F , we have

$$(243) \quad \{F, H\} = X_H(F) = \sum_i^m h^i(x) \frac{\partial F}{\partial x^i} .$$

Taking x^i as the function F , we get: $\{x^i, H\} = X_H(x^i) = h^i(x)$. So eq. 243 becomes

$$(244) \quad \{F, H\} = \sum_i^m \{x^i, H\} \frac{\partial F}{\partial x^i} .$$

If we now put x^i for H and H for F in eq. 244, we get

$$(245) \quad \{x^i, H\} = -\{H, x^i\} = -X_{x^i}(H) = -\sum_j^m \{x^j, x^i\} \frac{\partial H}{\partial x^j}.$$

Combining eq.s 244 and 245, we get the basic formula for the Poisson bracket of any two functions in terms of the Poisson bracket of local coordinates:

$$(246) \quad \{F, H\} = \sum_i^m \sum_j^m \{x^i, x^j\} \frac{\partial F}{\partial x^i} \frac{\partial H}{\partial x^j}.$$

We assemble these basic brackets, which we call the *structure functions* of the Poisson manifold,

$$(247) \quad J^{ij}(x) := \{x^i, x^j\} \quad i, j = 1, \dots, m$$

into a $m \times m$ anti-symmetric matrix of functions, $J(x)$, called the *structure matrix* of M . More precisely, it is the structure matrix for M relative to our coordinate system \mathbf{x} . Of course, the transformation of J under a coordinate change $x'^i := x'^i(x^1, \dots, x^m)$ is determined by setting $F := x'^i, H := x'^j$ in the basic formula eq. 246.

Then, writing ∇H for the (column) gradient vector of H , eq. 246 becomes

$$(248) \quad \{F, H\} = \nabla F \cdot J \nabla H.$$

For example, the canonical bracket on \mathbb{R}^{2n+l} , eq.236, written in the (q, p, z) coordinates, has the simple form

$$(249) \quad J = \begin{pmatrix} 0 & I & 0 \\ -I & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

where I is the $n \times n$ identity matrix.

We can write the Hamiltonian vector field, and the Hamilton's equations, associated with the function H in terms of J . Since

$$(250) \quad \{x^i, H\} = \sum_j^m \{x^i, x^j\} \frac{\partial H}{\partial x^j}$$

we get:

$$(251) \quad X_H = \sum_i^m \left(\sum_j^m J^{ij}(x) \frac{\partial H}{\partial x^j} \frac{\partial}{\partial x^i} \right),$$

or in matrix notation: $X_H = (J \nabla H) \cdot \partial_x$. Similarly, Hamilton's equations

$$(252) \quad \frac{dx^i}{dt} = \{x^i, H\}$$

get the matrix form

$$(253) \quad \frac{dx}{dt} = J(x) \nabla H(x) \quad ; \quad \text{i.e.} \quad \frac{dx^i}{dt} = \sum_j^m J^{ij}(x) \frac{\partial H}{\partial x^j}.$$

To summarize how we have generalized from the usual form of Hamilton's equations: compare eq. 253, 248 and 249 respectively with eq. 12, 18 and 3.

Note that not every $m \times m$ anti-symmetric matrix of functions on an m -dimensional manifold (or even: on an open subset of \mathbb{R}^m) is the structure matrix of a Poisson manifold: for the Jacobi identity constrains the functions. In fact it is readily shown that the Jacobi identity corresponds to the following m^3 partial differential equations governing the $J^{ij}(x)$, which are in general non-linear. Writing as usual ∂_l for $\partial/\partial x^l$:

$$(254) \quad \sum_{l=1}^m (J^{il} \partial_l J^{jk} + J^{kl} \partial_l J^{ij} + J^{jl} \partial_l J^{ki}) = 0 \quad i, j, k, = 1, \dots, m; \forall x \in M.$$

In particular, any constant anti-symmetric matrix J defines a Poisson structure.

5.2.4 The Poisson structure on \mathfrak{g}^*

We can now show that any m -dimensional Lie algebra \mathfrak{g} defines a Poisson structure, often called the *Lie-Poisson bracket*, on any m -dimensional vector space V . We proceed in two stages.

- (1) We first present the definition in a way that seems to depend on a choice of bases, both in \mathfrak{g} (where the definition makes a choice of structure constants) and in the space V .
- (2) Then we will see that choosing V to be \mathfrak{g}^* , the definition is in fact basis-independent.

This Poisson structure on \mathfrak{g}^* will be of central importance from now on. As Marsden and Ratiu write: 'Besides the Poisson structure on a symplectic manifold, the Lie-Poisson bracket on \mathfrak{g}^* , the dual of a Lie algebra, is perhaps the most fundamental example of a Poisson structure' (1999: 415). Here we return to our motivating discussion of Poisson manifolds, especially reasons (2) and (3) of Section 5.1: which concerned the rigid body and reduction, respectively. Indeed, we will see already in the Example at the end of this Subsection (Section 5.2.4.1) how the Lie-Poisson bracket on the special case $\mathfrak{g}^* := \mathfrak{so}(3)^*$ clarifies the theory of the rigid body. And we will see in Sections 7.2 and 7.3.3 how for any \mathfrak{g} , the Lie-Poisson bracket on \mathfrak{g}^* is induced by reduction, from the canonical Poisson (viz. symplectic) structure on the cotangent bundle T^*G . This will be our reduction theorem, that $T^*G/G \cong \mathfrak{g}^*$.

After (2), we will see that the Lie-Poisson bracket on \mathfrak{g}^* implies that Hamilton's equations on \mathfrak{g}^* can be expressed using ad^* : a form that will be needed later. This will be (3) below. Then we will turn in Section 5.2.4.1 to the example $\mathfrak{g}^* := \mathfrak{so}(3)^*$.

(1): A Poisson bracket on any vector space V :

Take a basis, say e_1, \dots, e_m , in \mathfrak{g} , and so structure constants c_{ij}^k (cf. eq. 52). Consider the space V as a manifold, and coordinatize it by taking a basis, $\epsilon_1, \dots, \epsilon_m$ say, determining coordinates x^1, \dots, x^m . We now define the Poisson bracket (in

this case, often called the *Lie-Poisson bracket*) between two smooth functions $F, H : V \rightarrow \mathbb{R}$ to be

$$(255) \quad \{F, H\} := \sum_{i,j,k=1}^m c_{ij}^k x^k \frac{\partial F}{\partial x^i} \frac{\partial H}{\partial x^j}.$$

This takes the form of eq. 246, with linear structure functions $J^{ij}(x) = \sum_k^m c_{ij}^k x^k$. One easily checks that anti-symmetry, and the Jacobi identity, for the structure constants, eq. 53, implies that these J^{ij} are anti-symmetric and obey their Jacobi identity eq. 254. So eq. 255 defines a Poisson bracket on V .

In particular, the associated Hamiltonian equations, eq.s 252 and 253, take the form

$$(256) \quad \frac{dx^i}{dt} = \sum_{j,k=1}^m c_{ij}^k x^k \frac{\partial H}{\partial x^j}.$$

(2): *The Lie-Poisson bracket on \mathfrak{g}^* :*

To give a basis-independent characterization of the Lie-Poisson bracket, we first recall that:

- (i) the gradient $\nabla F(x)$ of $F : V \rightarrow \mathbb{R}$ at any point $x \in V$ is in the dual space V^* of (continuous) linear functionals on V ;
- (ii) any finite-dimensional vector space is canonically, i.e. basis-independently, isomorphic to its double dual: $(V^*)^* \cong V$.

Then writing $\langle ; \rangle$ for the natural pairing between V and V^* , we have, for any $y \in V$

$$(257) \quad \langle \nabla F(x); y \rangle := \lim_{\tau \rightarrow 0} \frac{F(x + \tau y) - F(x)}{\tau}.$$

Now let us take V in our definition of the Lie-Poisson bracket to be \mathfrak{g}^* . So we will show that \mathfrak{g} makes \mathfrak{g}^* a Poisson manifold, in a basis-independent way. And let the basis $\epsilon_1, \dots, \epsilon_m$ be dual to the basis e_1, \dots, e_m of \mathfrak{g} . If $F : \mathfrak{g}^* \rightarrow \mathbb{R}$ is any smooth function, its gradient $\nabla F(x)$ at any point $x \in \mathfrak{g}^*$ is an element of $(\mathfrak{g}^*)^* \cong \mathfrak{g}$. One now checks that the Lie-Poisson bracket defined by eq. 255 has the basis-independent expression

$$(258) \quad \{F, H\}(x) = \langle x; [\nabla F(x), \nabla H(x)] \rangle, \quad x \in \mathfrak{g}^*$$

where $[\cdot, \cdot]$ is the ordinary Lie bracket on the Lie algebra \mathfrak{g} itself.

(3): *Hamilton's equations on \mathfrak{g}^* :*

We can also give a basis-independent expression of the Hamilton's equations eq. 256: viz. by expressing the Lie bracket in eq. 258 in terms of *ad*, as indicated by eq. 179.

Thus let $F \in \mathcal{F}(\mathfrak{g}^*)$ be an arbitrary smooth scalar function on \mathfrak{g}^* . By the chain rule

$$(259) \quad \frac{dF}{dt} = \mathbf{D}F(x) \cdot \dot{x} = \langle \dot{x}; \nabla F(x) \rangle .$$

But applying eq.s 179 and 198 to eq. 258 implies:

$$(260) \quad \begin{aligned} \{F, H\}(x) &= \langle x; [\nabla F(x), \nabla H(x)] \rangle = \\ &= \langle x; \text{ad}_{\nabla H(x)}(\nabla F(x)) \rangle = \langle \text{ad}_{\nabla H(x)}^*(x); \nabla F(x) \rangle . \end{aligned}$$

Since F is arbitrary and the pairing is non-degenerate, we deduce that Hamilton's equations take the form

$$(261) \quad \frac{dx}{dt} = \text{ad}_{\nabla H(x)}^*(x) .$$

5.2.4.1 Example: $\mathfrak{so}(3)$ and $\mathfrak{so}(3)^*$ As an example of the dual of a Lie algebra as a Poisson manifold, let us consider again our standard example $\mathfrak{so}(3)^*$. We will thereby make good our promise in (2) of Section 5.1, to show that Euler's equations for a rigid body are *already* in Hamiltonian form — in our generalized sense. We will also see why in the Chapter's second motto, Arnold mentions the three dual spaces, \mathbb{R}^{3*} , $\mathfrak{so}(3)^*$ and $T^*(SO(3))_g$; (cf. the discussion at the end of Section 3.4.4).

The Lie algebra $\mathfrak{so}(3)$ of $SO(3)$ has a basis e_1, e_2, e_3 representing infinitesimal rotations around the x -, y - and z -axes of \mathbb{R}^3 . As we have seen, we can think of these basis elements: as vectors in \mathbb{R}^3 with $[\cdot, \cdot]$ as elementary vector multiplication; or as anti-symmetric matrices with $[\cdot, \cdot]$ as the matrix commutator; or as left-invariant vector fields on $SO(3)$ with $[\cdot, \cdot]$ as the vector field commutator (i.e. Lie bracket).

Let $\epsilon_1, \epsilon_2, \epsilon_3$ be a dual basis for $\mathfrak{so}(3)^*$, with $x = x^1\epsilon_1 + x^2\epsilon_2 + x^3\epsilon_3$ a typical point therein. If $F : \mathfrak{so}(3)^* \rightarrow \mathbb{R}$, its gradient at x is the vector

$$(262) \quad \nabla F = \frac{\partial F}{\partial x^1}e_1 + \frac{\partial F}{\partial x^2}e_2 + \frac{\partial F}{\partial x^3}e_3 \in \mathfrak{so}(3).$$

Then eq. 258 tells us that, if we write $\mathfrak{so}(3)$ as \mathbb{R}^3 with \times for elementary vector multiplication, the Lie-Poisson bracket on $\mathfrak{so}(3)^*$ is

$$(263) \quad \{F, H\}(x) = x^1 \left(\frac{\partial F}{\partial x^3} \frac{\partial H}{\partial x^2} - \frac{\partial F}{\partial x^2} \frac{\partial H}{\partial x^3} \right) + \dots + x^3 \left(\frac{\partial F}{\partial x^2} \frac{\partial H}{\partial x^1} - \frac{\partial F}{\partial x^1} \frac{\partial H}{\partial x^2} \right)$$

$$(264) \quad = -x \cdot (\nabla F \times \nabla H).$$

So the structure matrix $J(x)$ is

$$(265) \quad J(x) = \begin{pmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{pmatrix}, \quad x \in \mathfrak{so}(3)^* .$$

Hamilton's equations corresponding to the Hamiltonian function $H(x)$ are therefore

$$(266) \quad \frac{dx}{dt} = x \times \nabla H(x) .$$

Now consider the Hamiltonian representing the kinetic energy of a free pivoted rigid body

$$(267) \quad H(x) = \frac{1}{2} \left(\frac{(x^1)^2}{I_1} + \frac{(x^2)^2}{I_2} + \frac{(x^3)^2}{I_2} \right),$$

in which the I_i are the moments of inertia about the three coordinate axes, and the x^i are the corresponding components of the *body* angular momentum. For this Hamiltonian, Hamilton's equations eq. 266 become

$$(268) \quad \frac{dx^1}{dt} = \frac{I_2 - I_3}{I_2 I_3} x^2 x^3, \quad \frac{dx^2}{dt} = \frac{I_3 - I_1}{I_3 I_1} x^3 x^1, \quad \frac{dx^3}{dt} = \frac{I_1 - I_2}{I_1 I_2} x^1 x^2,$$

Indeed, these are the Euler equations for a free pivoted rigid body. I shall not go into details about the rigid body. I only note that:

- (i) In the elementary theory of such a body, the magnitude L of the angular momentum is conserved, and eq. 268 describes the motion of the x^i on a sphere of radius L centred at the origin.
- (ii) In Section 5.4, we will return to seeing these spheres as the orbits of the co-adjoint representation of $SO(3)$ on $\mathfrak{so}(3)^*$ (cf. Section 4.5.2).
- (iii) Let us sum up this theme by saying, with Marsden and Ratiu (1999, p.11) that here we see: 'a simple and beautiful Hamiltonian structure for the rigid body equations'.

5.3 The symplectic foliation of Poisson manifolds

We first reformulate some ideas of Section 5.2 in more coordinate-independent language, starting with Section 5.2.3's idea of the structure matrix $J(x)$ (Section 5.3.1). Then we discuss canonical transformations on a Poisson manifold (Section 5.3.2). This will lead up to showing that any Poisson manifold is foliated by symplectic leaves (Section 5.3.3). Finally, we state a generalization of Darboux's theorem; and again take $\mathfrak{so}(3)$ as an example (Section 5.3.4).

5.3.1 The Poisson structure and its rank

We now pass from the structure matrix J , eq. 247, to a coordinate-independent object, the *Poisson structure* (also known as: co-symplectic structure), written B . Whereas J multiplied naive gradient vectors, as in eq. 248 and 253, B is to map the 1-form dH into its Hamiltonian vector field; as follows.

At each point x in a Poisson manifold M , there is a unique linear map B_x , which we will also write as B

$$(269) \quad B \equiv B_x : T_x^* M \rightarrow T_x M$$

such that

$$(270) \quad \mathbf{B}_x(dH(x)) = X_H(x).$$

For the requirement eq. 270 implies, by eq. 251, that for each $j = 1, \dots, m$

$$(271) \quad \mathbf{B}_x(dx^j) = \Sigma_i J^{ij}(x) \frac{\partial}{\partial x^i} \Big|_x$$

Since the differentials dx^i span T_x^*M , this fixes \mathbf{B}_x , by linearity. \mathbf{B}_x 's action on any one-form $\alpha = \Sigma a_j dx^j$ is:

$$(272) \quad \mathbf{B}_x(\alpha) = \Sigma_{i,j} J^{ij}(x) a_j \frac{\partial}{\partial x^i} \Big|_x$$

so that \mathbf{B}_x is essentially matrix multiplication by $J(x)$. Here, compare again eq. 252 and 253.

Here we recall that any linear map between (real finite-dimensional) vector spaces, $B : V \rightarrow W^*$, has an associated bilinear form B^\sharp on $V \times W^{**} \cong V \times W$ given by

$$(273) \quad B^\sharp(v, w) := \langle B(v); w \rangle .$$

Accordingly, some authors introduce the Poisson structure as a bilinear form $\mathbf{B}_x^\sharp : T_x^*M \times T_x^*M \rightarrow \mathbb{R}$, often called the *Poisson tensor*. Thus eq. 273 gives, for $\alpha, \beta \in T_x^*M$

$$(274) \quad \mathbf{B}_x^\sharp(\alpha, \beta) := \langle \mathbf{B}(\alpha), \beta \rangle .$$

\mathbf{B}_x^\sharp is antisymmetric, since the matrix $J(x)$ is. So, if we now let x vary over M , we can sum up in the traditional terminology of tensor analysis: \mathbf{B}^\sharp is an antisymmetric contravariant two-tensor field.

Example: Consider our first example, $M = \mathbb{R}^{2n+l}$ with the “usual bracket” eq. 236, from the start of Section 5.2.1. For any one-form

$$(275) \quad \alpha = \Sigma_{i=1}^n (a_i dq^i + b_i dp^i) + \Sigma_{j=1}^l c_j dz^j$$

we have

$$(276) \quad \mathbf{B}(\alpha) = \Sigma_{i=1}^n \left(b_i \frac{\partial}{\partial q^i} - a_i \frac{\partial}{\partial p^i} \right).$$

In this example the form of \mathbf{B} is the same from point to point. In particular, the kernel of \mathbf{B} has everywhere the same dimension, viz. l , the number of distinguished coordinates.

We now define the *rank* at x of a Poisson manifold M to be the rank of its Poisson structure \mathbf{B} at x , i.e. the dimension of the range of \mathbf{B}_x . This range is also the span of all the Hamiltonian vector fields on M at x :

$$(277) \quad \text{ran}(\mathbf{B}_x) := \{X \in T_x M : X = \mathbf{B}_x(\alpha), \text{ some } \alpha \in T_x^*M\} = \{X_H(x) : H : M \rightarrow \mathbb{R} \text{ smooth}\} .$$

So the rank of M at x is also equal to the dimension of B_x 's domain, i.e. $\dim(T_x^*M) = \dim(M)$, minus the dimension of the kernel, $\dim(B_x)$.

Since in local coordinates, B_x is given by multiplication by the structure matrix $J(x)$, the rank of M at x is the rank (the same in any coordinates) of the matrix $J(x)$. That $J(x)$ is anti-symmetric implies that the rank of M is *even*: cf. again the normal form of antisymmetric bilinear forms, eq. 2 and 3.

The manifold M being symplectic corresponds, of course, to the rank of B being everywhere maximal, i.e. equal to $\dim(M)$.

In this case, the kernel of B is trivial, and any distinguished function H is constant on M . For H is distinguished iff $X_H = 0$; and if the rank is maximal, then $dH = 0$, so that H is constant.

Besides, each of the Poisson structure and symplectic form on M determine the other. In particular, the Poisson tensor B^\sharp of eq. 274 is, up to a sign, the “contravariant cousin” of M 's symplectic form ω . For recall: (i) the relation between a symplectic manifold's Poisson bracket and its form, eq. 18, viz.

$$(278) \quad \{F, H\} = dF(X_H) = \omega(X_F, X_H) ;$$

and (ii) eq. 239 for Hamiltonian vector fields on a Poisson manifold, viz.

$$(279) \quad X_H(F) = \{F, H\} .$$

Applying these equations yields, if we start from eq. 274 and eq. 270:

$$(280) \quad B^\sharp(dH, dF) := \langle B(dH), dF \rangle = dF(X_H) = X_H(F) = \{F, H\} = \omega(X_F, X_H) .$$

We have also seen examples where the Poisson structure B is of non-maximal rank:

- (i) In our opening “canonical” example, the Poisson bracket eq. 236 on $M = \mathbb{R}^{2n+l}$ has rank $2n$ everywhere.
- (ii) In the Lie-Poisson structure on $\mathfrak{so}(3)^*$, the rank varies across the manifold: it is 2 everywhere, except at the origin $x = 0$ where it is 0. (Cf. the rank of the matrix J in eq. 265.)

5.3.2 Poisson maps

Already at the beginning of our development of Poisson manifolds, we saw that a scalar function $H : M \rightarrow \mathbb{R}$ defines equations of motion, with H as “Hamiltonian”, for all other functions $F : M \rightarrow \mathbb{R}$, of the familiar Poisson bracket type:

$$(281) \quad \dot{F} = \{F, H\} .$$

(Cf. Section 5.2.2, especially the remarks around eq. 239.) We now develop the generalization for Poisson manifolds of some related notions and results.

We say that a smooth map $f : M_1 \rightarrow M_2$ between Poisson manifolds $(M_1, \{\cdot, \cdot\}_1)$ and $(M_2, \{\cdot, \cdot\}_2)$ is *Poisson* or *canonical* iff it preserves the Poisson bracket. To be

precise: we first need the idea of the *pullback* of a function; cf. Section 3.1.2.1. In this context, the pullback f^* of a function $F : M_2 \rightarrow \mathbb{R}$ is given by

$$(282) \quad f^*F := F \circ f; \quad \text{i.e. } f^*F : x \in M_1 \mapsto F(f(x)) \in \mathbb{R}.$$

Then we say that $f : M_1 \rightarrow M_2$ is *Poisson* iff for all smooth functions $F, G : M_2 \rightarrow \mathbb{R}$ ($F, G \in \mathcal{F}(M_2)$)

$$(283) \quad f^*\{F, G\}_2 = \{f^*F, f^*G\}_1 ;$$

where by the definition eq. 282, the lhs $\equiv \{F, G\}_2 \circ f$, and the rhs $\equiv \{F \circ f, G \circ f\}_1$.

We note the special case where $M_1 = M_2 =: M$ and M is symplectic; i.e. the Poisson bracket is of maximal rank, and so defines a symplectic form on M , as in eq. 280. In this case, we return to the equivalence in Section 2.1.3's usual formulation of Hamiltonian mechanics, between preserving the Poisson bracket and preserving the symplectic form. That is: a map $f : M \rightarrow M$ on a symplectic manifold M is Poisson iff it is symplectic.

Besides, we already have for symplectic manifolds an infinitesimal version of the idea of a Poisson or symplectic map: viz. the idea of a locally Hamiltonian vector field; cf. Section 2.1.3. Similarly for Poisson manifolds, we will need the corresponding infinitesimal version of a Poisson map; but not till Section 6.1.1.

One can show (using in particular the Jacobi identity) that the flows of a Hamiltonian vector field are Poisson. (Here of course, $(M_1, \{, \}_1) = (M_2, \{, \}_2)$.) That is: if ϕ_τ is the flow of X_H (i.e. $\phi_\tau = \exp(\tau X_H)$), then

$$(284) \quad \phi_\tau^*\{F, G\} = \{\phi_\tau^*F, \phi_\tau^*G\} \quad \text{i.e. } \{F, G\} \circ \phi_\tau = \{F \circ \phi_\tau, G \circ \phi_\tau\} .$$

Similarly, one can readily show the equivalent proposition, that along the flow of a Hamiltonian vector field the Lie derivative of the Poisson tensor \mathbb{B}^\sharp vanishes. That is: for any smooth function $H : M \rightarrow \mathbb{R}$, we have:

$$(285) \quad \mathcal{L}_{X_H} \mathbb{B}^\sharp = 0 .$$

Since preserving the Poisson bracket implies in particular preserving its rank, it follows from eq. 284 (or from eq. 285) that:

If X_H is a Hamiltonian vector field on a Poisson manifold M , then for any $\tau \in \mathbb{R}$ and $x \in M$, the rank of M at $\exp(\tau X_H)(x)$ is the same as the rank at x . In other words: Hamiltonian vector fields are rank-invariant in the sense used in the general form of Frobenius' theorem (Section 3.3.2).

This result will be important for the foliation theorem for Poisson manifolds.

We will also need the result (also readily shown) that Poisson maps push Hamiltonian flows forward to Hamiltonian flows. More precisely: let $f : M_1 \rightarrow M_2$ be a Poisson map; so that at each $x \in M_1$, we have the derivative map on the tangent space, $Tf : (TM_1)_x \rightarrow (TM_2)_{f(x)}$. And let $H : M_2 \rightarrow \mathbb{R}$ be a smooth function. If ϕ_τ is the flow of X_H and ψ_τ is the flow (on M_1) of $X_{H \circ f}$, then:

$$(286) \quad \phi_\tau \circ f = f \circ \psi_\tau \quad \text{and} \quad Tf \circ X_{H \circ f} = X_H \circ f .$$

In particular, this square commutes:

$$(287) \quad \begin{array}{ccc} M_1 & \xrightarrow{f} & M_2 \\ \uparrow \psi_\tau & & \uparrow \phi_\tau \\ M_1 & \xrightarrow{f} & M_2 \end{array}$$

5.3.3 Poisson submanifolds: the foliation theorem

To state the foliation theorem for Poisson manifolds, we need the idea of a *Poisson immersion*, which leads to the closely related idea of a *Poisson submanifold*. In effect, these ideas combine the idea of a Poisson map with the ideas about injective immersions in (2) of Section 3.3.1. We recall from that discussion that for an injective immersion, $f : N \rightarrow M$, the range $f(N)$ is not necessarily a submanifold of M : but $f(N)$ is nevertheless called an ‘injectively immersed submanifold’ of M . (But as mentioned in Section 3.3.2, many treatments ignore this point: they in effect assume that an injective immersion f is also an embedding, i.e. a homeomorphism between N and $f(N)$, so that $f(N)$ is indeed a submanifold of M and f is a diffeomorphism.)

An injective immersion $f : N \rightarrow M$, with M a Poisson manifold, is called a *Poisson immersion* if any Hamiltonian vector field defined on an open subset of M containing $f(N)$ is in the range of the derivative map of f at $y \in N$, i.e. $\text{ran}(T_y f)$, at all points $f(y)$ for $y \in N$.

Being a Poisson immersion is equivalent to the following rather technical condition.

Characterization of Poisson immersions An injective immersion $f : N \rightarrow M$, with M a Poisson manifold, is a Poisson immersion iff: if $F, G : V \subset N \rightarrow \mathbb{R}$, where V is open in N , and if $\bar{F}, \bar{G} : U \rightarrow \mathbb{R}$ are extensions of $F \circ f^{-1}, G \circ f^{-1} : f(V) \rightarrow \mathbb{R}$ to an open neighbourhood U of $f(V)$ in M , then $\{\bar{F}, \bar{G}\} |_{f(V)}$ is well-defined and independent of the extensions.

The main point of this equivalence is that it ensures that if $f : N \rightarrow M$ is a Poisson immersion, then N has a Poisson structure, and $f : N \rightarrow M$ is a Poisson map. It is worth seeing how this comes about — by proving the equivalence.

Proof. Let $f : N \rightarrow M$ be a Poisson immersion, and let $F, G : V \subset N \rightarrow \mathbb{R}$ and let $\bar{F}, \bar{G} : U \supset f(V) \rightarrow \mathbb{R}$ be extensions of $F \circ f^{-1}, G \circ f^{-1} : f(V) \rightarrow \mathbb{R}$. Then for $y \in V$, there is a unique vector $v \in TN_y$ such that

$$(288) \quad X_{\bar{G}}(f(y)) = (T_y f)(v) .$$

So evaluating the Poisson bracket of \bar{F} and \bar{G} at $f(y)$ yields, by eq. 239,

$$(289) \quad \{\bar{F}, \bar{G}\}(f(y)) = d\bar{F}(f(y)) \cdot X_{\bar{G}}(f(y)) = d\bar{F}(f(y)) \cdot (T_y f)(v) = d(\bar{F} \circ f)(y) \cdot v \equiv dF(y) \cdot v .$$

So $\{\bar{F}, \bar{G}\}(f(y))$ is independent of the extension \bar{F} of $F \circ f^{-1}$. Since the Poisson bracket is antisymmetric, it is also independent of the extension \bar{G} of $G \circ f^{-1}$. So we can define a Poisson structure on N by defining for any y in an open $V \subset N$

$$(290) \quad \{F, G\}_N(y) := \{\bar{F}, \bar{G}\}_M(f(y)) .$$

This makes $f : N \rightarrow M$ a Poisson map, since for any \bar{F}, \bar{G} on M and any $y \in N$, we have that

$$(291) \quad [f^*\{\bar{F}, \bar{G}\}_M](y) \equiv \{[\bar{F}, \bar{G}]_M \circ f\}(y) = \{F, G\}_N(y) \equiv \{f^*\bar{F}, f^*\bar{G}\}_N(y) ;$$

where the middle equality uses eq. 290.

For the converse implication, assume that eq. 289 holds, and let $H : U \rightarrow \mathbb{R}$ be a Hamiltonian defined on an open subset U of M that intersects $f(N)$. Then as we have just seen, N is a Poisson manifold and $f : N \rightarrow M$ is a Poisson map. Because f is Poisson, it pushes $X_{H \circ f}$ to X_H . That is: eq. 286 implies that if $y \in N$ is such that $f(y) \in U$, then

$$(292) \quad X_H(f(y)) = (T_y f)(X_{H \circ f}(y)) .$$

So $X_H(f(y))$ is in the range of $T_y f$; so $f : N \rightarrow M$ is a Poisson immersion. ■

Now suppose that the inclusion $id : N \rightarrow M$ is a Poisson immersion. Then we call N a *Poisson submanifold* of M . We emphasise, in line with the warning we recalled from (2) of Section 3.3.1, that N need not be a submanifold of M ; but it is nevertheless called an ‘injectively immersed submanifold’ of M .

From the definition of a Poisson immersion, it follows that any Hamiltonian vector field must be tangent to a Poisson submanifold. In other words: writing \mathcal{X} for the system of Hamiltonian vector fields on M , and $\mathcal{X}|_x$ for their values at $x \in M$, we have: if N is a Poisson submanifold of M , and $x \in N$, $\mathcal{X}|_{x \subset TN_x}$.

For the special case where M is a symplectic manifold, we have $\mathcal{X}|_{x= T_x M}$, and the only Poisson submanifolds of M are its open sets.

Finally, we define the following equivalence relation on a Poisson manifold M . Two points $x_1, x_2 \in M$ are *on the same symplectic leaf* if there is a piecewise smooth curve in M joining them, each segment of which is an integral curve of a locally defined Hamiltonian vector field. An equivalence class of this equivalence relation is a *symplectic leaf*.

We can now state and prove that Poisson manifolds are foliated.

5.3.3.1 Foliation theorem for Poisson manifolds

The result is:

A Poisson manifold M is the disjoint union of its symplectic leaves. Each symplectic leaf is an injectively immersed Poisson submanifold, and the induced Poisson structure on the leaf is symplectic. The leaf through the point x , N_x say, has dimension equal to the rank of the Poisson structure at x ; and the tangent space to the leaf at x equals

$$(293) \quad TN_x = \text{ran}(\mathbf{B}_x) := \{X \in T_x M : X = \mathbf{B}_x(\alpha), \text{ some } \alpha \in T_x^* M\} \\ = \{X_H(x) : H \in \mathcal{F}(U), U \text{ open in } M \}$$

Proof. We apply the general form of Frobenius' theorem (Section 3.3.2) to the system \mathcal{X} of Hamiltonian vector fields on M . We know from eq. 242 (Section 5.2.2) that \mathcal{X} is involutive, and from eq. 284 above that it is rank-invariant. So by Frobenius' theorem, \mathcal{X} is integrable. The integral submanifolds are by definition given by the rhs of eq. 293. ■

One also readily shows that:

(i) One can evaluate the Poisson bracket of $F, G : M \rightarrow \mathbb{R}$ at $x \in M$ by restricting F and G to the symplectic leaf N_x through x , and evaluating the Poisson bracket that is defined by the symplectic form on the leaf N_x ; (i.e. the Poisson bracket defined in eq. 18).

(ii) A distinguished function is constant on any symplectic leaf N_x of M .

We end with two remarks. The first is a mathematical warning; the second concerns physical interpretation.

- (1) Recall our warning that symplectic leaves need not be submanifolds. This also means that all the distinguished functions being constants does *not* imply that the Poisson structure is non-degenerate. Indeed, one can readily construct an example in which the symplectic leaves are not manifolds, all distinguished functions are constants, and the Poisson structure is degenerate. Namely, one adapts an example mentioned before, in Section 3.4.3: the flows on the torus \mathbb{T}^2 that wind densely around it. (For more details about this example, cf. [Arnold, 1973, 160-167; 1989, 72-74] or [Butterfield, 2004a, Section 2.1.3.B]; for how to adapt it, cf. [Marsden and Ratiu, 1999, 347].
- (2) As we have seen, any integral curve of any Hamiltonian vector field X_H is confined to one of the symplectic leaves. So if we are interested only in the behaviour of a single solution through a point $x \in M$, we can restrict our attention to the symplectic leaf N_x through x : for the solution will always remain in N_x . But as stressed in Section 5.1, there are at least three good reasons not to ignore the more general Poisson structure!

5.3.4 Darboux's theorem

At the end of Section 2.1.1, we mentioned Darboux's theorem: it said that any symplectic manifold "looks locally like" a cotangent bundle. The generalization for Poisson manifolds says that any Poisson manifold "looks locally like" our canonical example on \mathbb{R}^m , $m = 2n + l$, given at the start of Section 5.2.1. More precisely, we have:

Let M be an m -dimensional Poisson manifold, and let $x \in M$ be a point with an open neighbourhood $U \subset M$ throughout which the rank is a constant $2n \leq m$. Then defining $l := m - 2n$, there is a possibly smaller neighbourhood $U' \subset U$ of x , on which there exist local coordinates $(q, p, z) = (q^1, \dots, q^n, p^1, \dots, p^n, z^1, \dots, z^l)$, for which the Poisson bracket takes the form

$$(294) \quad \{F, H\} := \sum_i^n \left(\frac{\partial F}{\partial q^i} \frac{\partial H}{\partial p^i} - \frac{\partial F}{\partial p^i} \frac{\partial H}{\partial q^i} \right).$$

(So the Poisson brackets for the coordinate functions take the now-familiar form given by eq. 237 and 238.) The symplectic leaves of M intersect the coordinate chart in the slices $\{z^1 = c_1, \dots, z^l = c_l\}$ given by constant values of the distinguished coordinates z .

We shall not give the proof. Suffice it to say that:

- (i) Like Darboux's theorem for symplectic manifolds: it proceeds by induction on the "half-rank" n ; and it begins by taking any function F as the "momentum" p^1 and constructing the canonically conjugate coordinate q^1 such that $\{q^1, p^1\} = 1$.
- (ii) The induction step invokes a version of Frobenius' theorem in which the fact that the rank $2n$ is constant throughout U secures a coordinate system in which the $2n$ -dimensional integral manifolds are given by slices defined by constant values of the remaining l coordinates. The Poisson structure then secures that these remaining coordinates are distinguished.

5.3.4.1 Example: $\mathfrak{so}(3)^*$ yet again We illustrate (1) the foliation theorem and (2) Darboux's theorem, with $\mathfrak{so}(3)^*$; whose Lie-Poisson structure we described in Section 5.2.4.1.

- (1) At $x \in \mathfrak{so}(3)^*$, the subspace $\mathcal{X}|_x := \{X_H(x) : H \in \mathcal{F}(U), U \text{ open in } M\}$ of values of locally Hamiltonian vector fields is spanned by $e_1 := y\partial_z - z\partial_y$ representing infinitesimal rotation about the x -axis (cf. eqns. 48, 107 and 154); $e_2 := z\partial_x - x\partial_z$ for rotation about the y -axis; and $e_3 := x\partial_y - y\partial_x$ for rotation about the z -axis. If $x \neq 0$, these vectors span a two-dimensional subspace of $T\mathfrak{so}(3)^*$; viz. the tangent plane to the sphere $S_{|x|}$ of radius $|x|$ centred at the origin. So the foliation theorem implies that $\mathfrak{so}(3)^*$'s symplectic leaves are these spheres; and the origin.

We can compute the Poisson bracket of $F, G : S_{|x|} \rightarrow \mathbb{R}$ by extending F and G to a neighbourhood of $S_{|x|}$; cf. eq. 290. That is: we can consider extensions $\bar{F}, \bar{G} : U \supset S_{|x|} \rightarrow \mathbb{R}$, and calculate the Poisson bracket in $\mathfrak{so}(3)^*$, whose Poisson structure we already computed in eq. 264.

Adopting spherical polar coordinates with $r = |x|$, i.e. $x^1 = r \cos \theta \sin \phi, x^2 = r \sin \theta \sin \phi, x^3 = r \cos \phi$, we can define \bar{F}, \bar{G} merely by $\bar{F}(r, \theta, \phi) := F(\theta, \phi)$,

$\bar{G}(r, \theta, \phi) := G(\theta, \phi)$; so that the partial derivatives with respect to the spherical angles θ, ϕ are equal, i.e. $\bar{F}_\theta = F_\theta, \bar{F}_\phi = F_\phi, \bar{G}_\theta = G_\theta, \bar{G}_\phi = G_\phi$.

Besides, eq. 246 implies that we need only calculate the Poisson bracket in $\mathfrak{so}(3)^*$ of the spherical angles θ and ϕ . So eq. 264 gives

$$(295) \quad \{\theta, \phi\} = -x \cdot (\nabla\theta \times \nabla\phi) = \frac{-1}{r \sin \phi};$$

and eq. 290 and 246 give

$$(296) \quad \{F, G\} = \{\bar{F}, \bar{G}\} = \frac{-1}{r \sin \phi} (F_\theta G_\phi - F_\phi G_\theta).$$

- (2) $z := x^3$ defines the Hamiltonian vector field $X_z = x^2 \partial_{x^1} - x^1 \partial_{x^2}$ that generates clockwise rotation about the $z \equiv x^3$ -axis. So away from the origin the polar angle $\theta := \arctan(x^2/x^1)$ has a Poisson bracket with z equal to: $\{\theta, z\} = X_z(\theta) = -1$. Expressing $F, H : \mathfrak{so}(3)^* \rightarrow \mathbb{R}$ in terms of the coordinates z, θ and $r := |x|$, we find that the Lie-Poisson bracket is: $\{F, H\} = F_z H_\theta - F_\theta H_z$. So (z, θ, r) are canonical coordinates.

5.4 The symplectic structure of the co-adjoint representation

Section 5.2.4 described how the dual \mathfrak{g}^* of a finite-dimensional Lie algebra of a Lie group G has the structure of a Poisson manifold. In this case, the foliation established in the previous Subsection has an especially neat interpretation. Namely: the leaves are the orbits of the co-adjoint representation of G on \mathfrak{g}^* .

This symplectic structure in the co-adjoint representation sums up themes from Sections 4.5 (especially 4.5.2), and 5.2.4 and 5.3. In particular, it connects two properties of the Lie bracket in \mathfrak{g} , which we have already seen: viz.

- (i) The Lie bracket in \mathfrak{g} gives the infinitesimal generators of the adjoint action; cf. eq. 179.
- (ii) The Lie bracket in \mathfrak{g} defines (in a basis-independent way) a Lie-Poisson bracket on \mathfrak{g}^* , thus making \mathfrak{g}^* a Poisson manifold. (Cf. the definition in eq. 255, shown to be basis-independent by eq. 258.)

In fact, there is a wealth of instructive results and examples about the structure of the co-adjoint representation: we will only scratch the surface — as in other Sections! We will give a proof, under a simplifying assumption, of one main result; and then make a few remarks about other results.

The result is:

The orbits of the co-adjoint representation are \mathfrak{g}^* 's leaves

Let G be a Lie group, with its co-adjoint representation Ad^* on \mathfrak{g}^* .

That is, recalling eq. 193, we have:

$$(297) \quad Ad^* : G \rightarrow \text{End}(\mathfrak{g}^*), \quad Ad^*_{g^{-1}} = (T_e(R_g \circ L_{g^{-1}}))^*.$$

The orbits of this representation are the symplectic leaves of \mathfrak{g}^* , taken as equipped with its natural Poisson structure, i.e. the Lie-Poisson bracket eq. 258.

Proof. We shall prove this under the simplifying assumption that the co-adjoint action of G on \mathfrak{g}^* is proper. (We recall from the definition of proper actions, eq. 147, that for any compact Lie group, such as $SO(3)$, this condition is automatically satisfied.) Then we know from result (3) and eq. 160, at the end of Section 4.4, that this implies that the co-adjoint orbit $\text{Orb}(\alpha)$ of any $\alpha \in \mathfrak{g}^*$ is a closed submanifold of \mathfrak{g}^* , and that the tangent space to $\text{Orb}(\alpha)$ at a point $\beta \in \text{Orb}(\alpha)$ is

$$(298) \quad T\text{Orb}(\alpha)_\beta = \{ \xi_{\mathfrak{g}^*}(\beta) : \xi \in \mathfrak{g} \}.$$

We will see shortly how this assumption implies that \mathfrak{g}^* 's symplectic leaves are submanifolds.²¹

We now argue as follows. For $\xi \in \mathfrak{g}$, consider the scalar function on \mathfrak{g}^* , $K_\xi : \alpha \in \mathfrak{g}^* \mapsto K_\xi(\alpha) := \langle \alpha; \xi \rangle \in \mathbb{R}$; and its Hamiltonian vector field X_{K_ξ} . At each $\alpha \in \mathfrak{g}^*$, the gradient $\nabla K_\xi(\alpha) \equiv dK_\xi(\alpha)$, considered as an element of $(T^*\mathfrak{g}^*)_\alpha \cong \mathfrak{g}$, is just ξ itself. Now we will compute $X_{K_\xi}(F)(\alpha)$ for any $F : \mathfrak{g}^* \rightarrow \mathbb{R}$ and any $\alpha \in \mathfrak{g}^*$, using in order:

- (i) the intrinsic definition of the Lie-Poisson bracket on \mathfrak{g}^* , eq. 258;
- (ii) the fact that the infinitesimal generator of the adjoint action is the Lie bracket in \mathfrak{g} , eq. 179;
- (iii) the fact that the derivative ad^* of the co-adjoint action Ad^* is, up to a sign, the adjoint of ad_ξ ; eq. 198.

Thus we get, for all $F : \mathfrak{g}^* \rightarrow \mathbb{R}$ and $\alpha \in \mathfrak{g}^*$:

$$(299) \quad X_{K_\xi}(F)(\alpha) \equiv \{F, K_\xi\}(\alpha) = \langle \alpha; [\nabla F(\alpha), \nabla K_\xi(\alpha)] \rangle$$

$$(300) \quad = \langle \alpha; [\nabla F(\alpha), \xi] \rangle = - \langle \alpha; [\xi, \nabla F(\alpha)] \rangle$$

$$(301) \quad = - \langle \alpha; ad_\xi(\nabla F(\alpha)) \rangle$$

$$(302) \quad = \langle ad_\xi^*(\alpha); \nabla F(\alpha) \rangle.$$

But on the other hand, the vector field X_{K_ξ} is uniquely determined by its action on all such functions F at all $\alpha \in \mathfrak{g}^*$:

$$(303) \quad X_{K_\xi}(F)(\alpha) \equiv \langle X_{K_\xi}(\alpha); \nabla F(\alpha) \rangle.$$

So we conclude that at each $\alpha \in \mathfrak{g}^*$:

$$(304) \quad X_{K_\xi} = ad_\xi^*.$$

²¹To verify that our condition is indeed simplifying — i.e. that in general the co-adjoint orbits in \mathfrak{g}^* are not submanifolds — consider the example in Marsden and Ratiu (1999: 14.1.(f), p. 449); taken from Kirillov (1976: 293).

But the subspace $\mathcal{X}|_\alpha$ of values at α of Hamiltonian vector fields is spanned by the $X_{K_\xi}(\alpha)$, with ξ varying through \mathfrak{g} . And as ξ varies through \mathfrak{g} , $ad_\xi^*(\alpha)$ is the tangent space $T\text{Orb}(\alpha)_\alpha$ to the co-adjoint orbit $\text{Orb}(\alpha)$ of G through α . So

$$(305) \quad \mathcal{X}|_\alpha = T\text{Orb}(\alpha)_\alpha .$$

So the integral submanifolds of the system \mathcal{X} of Hamiltonian vector fields, which are the symplectic leaves of \mathfrak{g}^* by Section 5.3.3.1's foliation theorem, are the co-adjoint orbits. ■

For the illustration of this theorem by our standard example, $\mathfrak{so}(3)^*$, cf. our previous discussions of it: in Section 4.5.2 for its co-adjoint structure; in Section 5.2.4.1 for its Lie-Poisson structure; and in Section 5.3.4.1 for its symplectic leaf structure.

We end this Subsection by stating two other results. They are not needed later, but they are enticing hints of how rich is the theory of co-adjoint orbits.

- (1) For each $g \in G$, the co-adjoint map $Ad_g^* : \mathfrak{g}^* \rightarrow \mathfrak{g}^*$ is a Poisson map that preserves the symplectic leaves of \mathfrak{g}^* .
- (2) A close cousin of the theorem just proven is that the Lie bracket on \mathfrak{g} defines (via its definition of the Lie-Poisson bracket on \mathfrak{g}^* , eq. 258) a symplectic form, i.e. a non-degenerate closed two-form, on each co-adjoint orbit, by:

$$(306) \quad \omega(\alpha)(ad_\xi^*(\alpha), ad_\eta^*(\alpha)) := \langle \alpha ; [\xi, \eta]_{\mathfrak{g}} \rangle , \quad \forall \alpha \in \mathfrak{g}^* , \forall \xi, \eta \in \mathfrak{g} .$$

This theorem is proven in detail (without our simplifying assumption that G 's action is proper) by Marsden and Ratiu [1999, Thm 14.3.1, pp. 453–456]; and much more briefly by Arnold [1989, 321, 376–377, 457]; and rather differently (even without using the notion of a Poisson manifold!) in [Abraham and Marsden, 1978, 302–303].

5.5 Quotients of Poisson manifolds

We now end Section 5 with the simplest general theorem about quotienting a Lie group action on a Poisson manifold, so as to get a quotient space (set of orbits) that is itself a Poisson manifold. So this theorem combines themes from Sections 4 — in particular, the idea from Section 4.3.0.2 that for a free and proper group action, the orbits and quotient space are manifolds — with material about Poisson manifolds from Section 5.2. (The material in Sections 5.3 and 5.4 will not be needed.) This theorem will be important in Section 7. We call this result the

Poisson reduction theorem: Suppose the Lie group G acts on Poisson manifold M in such a way that each $\Phi_g : M \rightarrow M$ is a Poisson map. Suppose also that the quotient space M/G is a manifold and the projection $\pi : M \rightarrow M/G$ is a smooth submersion (say because G 's action on M is free and proper, cf. Section 4.3.0.3). Then there is a unique Poisson structure on M/G such that π is a Poisson map. The Poisson bracket on M/G is called the *reduced* Poisson bracket.

Proof. Let us first assume that M/G is a Poisson manifold and that π is a Poisson map; and show uniqueness. We first note that for any $f : M/G \rightarrow \mathbb{R}$, the function $\bar{f} := f \circ \pi : M \rightarrow \mathbb{R}$ is obviously the unique G -invariant function on M that projects by π to f . That is: if $[x] \equiv \text{Orb}(x) \equiv G \cdot x$ is the orbit of $x \in M$, then \bar{f} assigns the same value $f([x])$ to all elements of the orbit $[x]$. Besides, in terms of pullbacks (eq. 282), $\bar{f} = \pi^* f$.

Then the condition that π be Poisson, eq. 283, is that for any two smooth scalars $f, h : M/G \rightarrow \mathbb{R}$, we have an equation of smooth scalars on M :

$$(307) \quad \{f, h\}_{M/G} \circ \pi = \{f \circ \pi, h \circ \pi\}_M = \{\bar{f}, \bar{h}\}_M$$

where the subscripts indicate on which space the Poisson bracket is defined. Since π is surjective, eq. 307 determines the value $\{f, h\}_{M/G}$ uniquely.

But eq. 307 also *defines* $\{f, h\}_{M/G}$ as a Poisson bracket; in two stages. (1): The facts that Φ_g is Poisson, and \bar{f} and \bar{h} are constant on orbits imply that

$$(308) \quad \{\bar{f}, \bar{h}\}(g \cdot x) = (\{\bar{f}, \bar{h}\} \circ \Phi_g)(x) = \{\bar{f} \circ \Phi_g, \bar{h} \circ \Phi_g\}(x) = \{\bar{f}, \bar{h}\}(x).$$

That is: $\{\bar{f}, \bar{h}\}$ is also constant on orbits, and so defines $\{f, h\}$ uniquely.

(2): We show that $\{f, h\}$, as thus defined, is a Poisson structure on M/G , by checking that the required properties, such as the Jacobi identity, follow from the Poisson structure $\{, \}_M$ on M . ■

This theorem is a “prototype” for material to come. We spell this out in two brief remarks, which look forward to the following two Sections.

(1): *Other theorems:*

This theorem is one of many that yield new Poisson manifolds and symplectic manifolds from old ones by quotienting. In particular, as we will see in detail in Section 7, this theorem is exemplified by the case where $M = T^*G$ (so here M is symplectic, since it is a cotangent bundle), and G acts on itself by left translations, and so acts on T^*G by a cotangent lift. In this case, we will have $M/G \cong \mathfrak{g}^*$; and the reduced Poisson bracket just defined, by eq. 307, will be the Lie-Poisson bracket we have already met in Section 5.2.4.

(2): *Reduction of dynamics:*

Using this theorem, we can already fill out a little what is involved in reduced dynamics; which we only glimpsed in our introductory discussions, in Section 2.3 and 5.1. We can make two basic points, as follows.

- (A) If H is a G -invariant Hamiltonian function on M , it defines a corresponding function h on M/G by $H = h \circ \pi$. The fact that Poisson maps push Hamiltonian flows forward to Hamiltonian flows (eq. 286) implies, since π is Poisson, that π transforms X_H on M to X_h on M/G . That is:

$$(309) \quad T\pi \circ X_H = X_h \circ \pi \quad ;$$

i.e. X_H and X_h are π -related. Accordingly, we say that the Hamiltonian system X_H on M *reduces* to that on M/G .

- (B) We shall see in Section 6.2 that G -invariance of H is associated with a family of conserved quantities (constants of the motion, first integrals), viz. a constant of the motion $J(\xi) : M \rightarrow \mathbb{R}$ for each $\xi \in \mathfrak{g}$. Here, J being conserved means $\{J, H\} = 0$; just as in our discussion of Noether's theorem in ordinary Hamiltonian mechanics (Section 2.1.3). Besides, if J is also G -invariant, then the corresponding function j on M/G is conserved by X_h since

$$(310) \quad \{j, h\} \circ \pi = \{J, H\} = 0 \text{ implies } \{j, h\} = 0.$$

6 SYMMETRY AND CONSERVATION REVISITED: MOMENTUM MAPS

We now develop the topics of symmetry and conserved quantities (and so Noether's theorem) in the context of Poisson manifolds. At the centre of these topics lies the idea of a *momentum map* of a Lie group action on a Poisson manifold; which we introduce in Section 6.1. This is the modern geometric generalization of a conserved quantity, such as linear or angular momentum for the Euclidean group — hence the name. Formally, it will be a map \mathbf{J} from the Poisson manifold M to the dual \mathfrak{g}^* of the Lie algebra of the symmetry group G . Since its values lie in a vector space, it has components. So our description of conserved quantities will no longer be “one-dimensional”, i.e. focussed on a single vector field in the state space, as it was in Section 2.1.3. The map \mathbf{J} will be associated with a linear map J from \mathfrak{g} to $\mathcal{F}(M)$, the scalar functions on the manifold M . That is: for each $\xi \in \mathfrak{g}$, $J(\xi)$ will be a conserved quantity if the Hamiltonian H is invariant under the infinitesimal generator ξ_M , i.e. if $\xi_M(H) = 0$.

The conservation of momentum maps will be expressed by the Poisson manifold version of Noether's theorem (Section 6.2), and illustrated by the familiar examples of linear and angular momentum (Section 6.3). Then we discuss the equivariance of momentum maps, with respect to the co-adjoint representation of G on \mathfrak{g}^* ; Section 6.4. Finally in Section 6.5, we discuss the crucial special case of momentum maps on cotangent bundles, again with examples.

6.1 Canonical actions and momentum maps

We first apply the definition of Poisson maps (from Section 5.3.2) to group actions (Section 6.1.1). This will lead to the idea of the momentum map (Section 6.1.2).

6.1.1 Canonical actions and infinitesimal generators

Let G be a Lie group acting on a Poisson manifold M by a smooth left action $\Phi : G \times M \rightarrow M$; so that as usual we write $\Phi_g : x \in M \mapsto \Phi_g(x) := g \cdot x \in M$. As in the definition of a Poisson map (eq. 283), we say the action is *canonical* if

$$(311) \quad \Phi_g^*\{F_1, F_2\} = \{\Phi_g^*F_1, \Phi_g^*F_2\}$$

for any $F_1, F_2 \in \mathcal{F}(M)$ and any $g \in G$. If M is symplectic with symplectic form ω , then the action is canonical iff it is symplectic, i.e. $\Phi_g^*\omega = \omega$ for all $g \in G$.

We will be especially interested in the infinitesimal version of this notion; and so with infinitesimal generators of actions. We recall from eq. 152 that the infinitesimal generator of the action corresponding to a Lie algebra element $\xi \in \mathfrak{g}$ is the vector field ξ_M on M obtained by differentiating the action with respect to g at the identity in the direction ξ :

$$(312) \quad \xi_M(x) = \frac{d}{d\tau}[\exp(\tau\xi) \cdot x] \Big|_{\tau=0}.$$

So we differentiate eq. 311 with respect to g in the direction ξ , to give:

$$(313) \quad \xi_M(\{F_1, F_2\}) = \{\xi_M(F_1), F_2\} + \{F_1, \xi_M(F_2)\}.$$

Such a vector field ξ_M is called an *infinitesimal Poisson automorphism*.

Side-remark: We will shortly see that it is the universal quantification over $g \in G$ in eq. 311, and correspondingly in eq. 313 and 315 below, that means our description of conserved quantities is no longer focussed on a single vector field; and in particular, that a momentum map representing a conserved quantity has components.

In the symplectic case, differentiating $\Phi_g^*\omega = \omega$ implies that the Lie derivative $\mathcal{L}_{\xi_M}\omega$ of ω with respect to ξ vanishes: $\mathcal{L}_{\xi_M}\omega = 0$. We saw in Section 2.1.3 that this is equivalent to ξ_M being locally Hamiltonian, i.e. there being a local scalar $J : U \subset M \rightarrow \mathbb{R}$ such that $\xi_M = X_J$. This was how Section 2.1.3 vindicated eq. 19's “one-liner” approach to Noether’s theorem: because the vector field X_f is locally Hamiltonian, it preserves the symplectic structure, i.e. Lie-derives the symplectic form $\mathcal{L}_{X_f}\omega = 0$ — as a symmetry should.

We also saw in result (2) at the end of Section 3.2.2 that the “meshing”, up to a sign, of the Poisson bracket on scalars with the Lie bracket on vector fields implied that the locally Hamiltonian vector fields form a Lie subalgebra of the Lie algebra $\mathcal{X}(M)$ of all vector fields.

Turning to the context of Poisson manifolds, we need to note two points. The first is a similarity with the symplectic case; the second is a contrast.

- (1) One readily checks, just by applying eq. 313, that the infinitesimal Poisson automorphisms are closed under the Lie bracket. So we write the Lie algebra of these vector fields as $\mathcal{P}(M)$: $\mathcal{P}(M) \subset \mathcal{X}(M)$.
- (2) On the other hand, Section 2.1.3’s equivalence between a vector field being locally Hamiltonian and preserving the geometric structure of the state-space breaks down.

Agreed, the first implies the second: a locally Hamiltonian vector field preserves the Poisson bracket. We noted this already in Section 5.3.2. The differential

statement was that such a field X_H Lie-derives the Poisson tensor: $\mathcal{L}_{X_H} \mathbf{B}^\sharp = 0$ (eq. 285). The finite statement was that the flows of such a field are Poisson maps: $\phi_\tau^* \{F, G\} = \{\phi_\tau^* F, \phi_\tau^* G\}$ (eq. 284).

But the converse implication fails: an infinitesimal Poisson automorphism on a Poisson manifold need *not* be locally Hamiltonian. For example, make \mathbb{R}^2 a Poisson manifold by defining the Poisson structure

$$(314) \quad \{F, H\} = x \left(\frac{\partial F}{\partial x} \frac{\partial H}{\partial y} - \frac{\partial H}{\partial x} \frac{\partial F}{\partial y} \right) ;$$

then the vector field $X = \partial/\partial y$ in a neighbourhood of a point on the y -axis is a non-Hamiltonian infinitesimal Poisson automorphism.

This point will affect the formulation of Noether's theorem for Poisson manifolds, in Section 6.2.

Nevertheless, we shall from now on be interested in cases where for all ξ , ξ_M is globally Hamiltonian. This means there is a map $J : \mathfrak{g} \rightarrow \mathcal{F}(M)$ such that

$$(315) \quad X_{J(\xi)} = \xi_M$$

for all $\xi \in \mathfrak{g}$. There are three points we need to note about this condition.

- (1) Since the right hand side of eq. 315 is linear in ξ , we can require such a J to be a linear map. For given any J obeying eq. 315, we can take a basis e_1, \dots, e_m of \mathfrak{g} and define a new linear \bar{J} by setting, for any $\xi = \xi^i e_i$, $\bar{J}(\xi) := \xi^i J(e_i)$.
- (2) Eq. 315 does not determine $J(\xi)$. For by the linearity of the map $B : dJ(\xi) \mapsto X_{J(\xi)}$, we can add to such a $J(\xi)$ any distinguished function, i.e. an $F : M \rightarrow \mathbb{R}$ such that $X_F = 0$. That is: $X_{J(\xi)+F} \equiv X_{J(\xi)}$. (Of course, in the symplectic case, the only distinguished functions are constants.)
- (3) It is worth expressing eq. 315 in terms of Poisson brackets. Recalling that for any $F, H \in \mathcal{F}(M)$, we have $X_H(F) = \{F, H\}$, this equation becomes

$$(316) \quad \{F, J(\xi)\} = \xi_M(F) \quad , \quad \forall F \in \mathcal{F}(M), \quad \forall \xi \in \mathfrak{g} .$$

We will also need the following result:

$$(317) \quad X_{J([\xi, \eta])} = X_{\{J(\xi), J(\eta)\}_M} .$$

To prove this, we just apply two previous results, each giving a Lie algebra anti-homomorphism.

- (i) Result (4) at the end of Section 4.4: for any left action of Lie group G on any manifold M , the map $\xi \mapsto \xi_M$ is a Lie algebra anti-homomorphism between \mathfrak{g} and the Lie algebra \mathcal{X}_M of all vector fields on M :

$$(318) \quad (a\xi + b\eta)_M = a\xi_M + b\eta_M \quad ; \quad [\xi_M, \eta_M] = -[\xi, \eta]_M \quad \forall \xi, \eta \in \mathfrak{g}, \text{ and } a, b \in \mathbb{R} .$$

(ii) The “meshing” up to a sign, just as in the symplectic case, of the Poisson bracket on scalars with the Lie bracket on vector fields, as in eq. 242 at the end of Section 5.2.2:

$$(319) \quad X_{\{F,H\}} = -[X_F, X_H] = [X_H, X_F] .$$

So for a Poisson manifold M , the map $F \in \mathcal{F}(M) \mapsto X_F \in \mathcal{X}(M)$ is a Lie algebra anti-homomorphism.

Applying (i) and (ii), we deduce eq. 317 by:

$$(320) \quad X_{J([\xi,\eta])} = [\xi, \eta]_M = -[\xi_M, \eta_M] = -[X_{J(\xi)}, X_{J(\eta)}] = X_{\{J(\xi), J(\eta)\}_M} .$$

6.1.2 Momentum maps introduced

So suppose that there is a canonical left action of G on a Poisson manifold M . And suppose there is a linear map $J : \mathfrak{g} \rightarrow \mathcal{F}(M)$ such that

$$(321) \quad X_{J(\xi)} = \xi_M$$

for all $\xi \in \mathfrak{g}$.

The two requirements — that the action be infinitesimally canonical (i.e. each $\xi_M \in \mathcal{P}(M)$) and that each ξ_M be globally Hamiltonian — can be expressed as requiring that there be a $J : \mathfrak{g} \rightarrow \mathcal{F}(M)$ such that there is a commutative diagram. Namely, the map $\xi \in \mathfrak{g} \mapsto \xi_M \in \mathcal{P}(M)$ is to equal the composed map:

$$(322) \quad \mathfrak{g} \xrightarrow{J} \mathcal{F}(M) \xrightarrow{F \mapsto X_F} \mathcal{P}(M) .$$

Then the map $\mathbf{J} : M \rightarrow \mathfrak{g}^*$ defined by

$$(323) \quad \langle \mathbf{J}(x) ; \xi \rangle := J(\xi)(x)$$

for all $\xi \in \mathfrak{g}$ and $x \in M$, is called the *momentum map* of the action.

Another way to state this definition is as follows. Any smooth function $\mathbf{J} : M \rightarrow \mathfrak{g}^*$ defines at each $\xi \in \mathfrak{g}$ a scalar $J(\xi) : x \in M \mapsto (\mathbf{J}(x))(\xi) \in \mathbb{R}$. By taking $J(\xi)$ as a Hamiltonian function, one defines a Hamiltonian vector field $X_{J(\xi)}$. But since G acts on M , each $\xi \in \mathfrak{g}$ defines a vector field on M , viz. ξ_M . So we say that \mathbf{J} is a *momentum map* for the action if for each $\xi \in \mathfrak{g}$, these two vector fields are identical: $X_{J(\xi)} = \xi_M$.

Three further remarks by way of illustrating this definition:

(1): *An isomorphism:*

One readily checks that eq. 323 defines an isomorphism between the space of smooth maps \mathbf{J} from M to \mathfrak{g}^* , and the space of linear maps J from \mathfrak{g} to scalar functions $\mathcal{F}(M)$. We can take J to define \mathbf{J} by saying that at each $x \in M$, $\mathbf{J}(x) : \xi \in \mathfrak{g} \mapsto \mathbf{J}(x)(\xi) \in \mathbb{R}$ is to be given by the composed map

$$(324) \quad \mathfrak{g} \xrightarrow{J} \mathcal{F}(M) \xrightarrow{|_x} \mathbb{R} ,$$

where $|_x$ means evaluation at $x \in M$. Or we can take \mathbf{J} to define J by saying that at each $\xi \in \mathfrak{g}$, $J(\xi) : x \in M \mapsto J(\xi)(x) \in \mathbb{R}$ is to be given by the composed map

$$(325) \quad M \xrightarrow{\mathbf{J}} \mathfrak{g}^* \xrightarrow{|\xi} \mathbb{R} ,$$

where $|\xi$ means evaluation at $\xi \in \mathfrak{g}$.

(2): *Differential equations for the momentum map:*

Using Hamilton's equations, we can readily express the definition of momentum map as a set of differential equations. Recall that on a Poisson manifold, Hamilton's equations are determined by eq. 270, which was that at each $x \in M$

$$(326) \quad \mathbf{B}_x(dH(x)) = X_H(x) ;$$

or in local coordinates $x^i, i = 1, \dots, m \equiv \dim(M)$, with $J^{ij}(x) \equiv \{x^i, x^j\}$ the structure matrix,

$$(327) \quad \mathbf{B}_x\left(\frac{\partial H}{\partial x^j} dx^j\right) = \Sigma_{i,j} J^{ij}(x) \frac{\partial H}{\partial x^j} \frac{\partial}{\partial x^i} \Big|_x ;$$

(cf. eq. 272). So in local coordinates, Hamilton's equations are given by eq. 253, which was:

$$(328) \quad \frac{dx^i}{dt} = \Sigma_j^m J^{ij}(x) \frac{\partial H}{\partial x^j} .$$

So the condition for a momentum map $X_{J(\xi)} = \xi_M$ is that for all $\xi \in \mathfrak{g}$ and all $x \in M$

$$(329) \quad \mathbf{B}_x(d(J(\xi))(x)) = \xi_M(x) .$$

In coordinates, this is the requirement that for all $i = 1, \dots, m$

$$(330) \quad \Sigma_j^m J^{ij}(x) \frac{\partial J(\xi)}{\partial x^j} = (\xi_M)^i(x) ,$$

where — apologies! — the two J s on the left hand side have very different meanings.

In the symplectic case, $\dim(M) \equiv m = 2n$ and we have Hamilton's equations as eq. 15, viz.

$$(331) \quad \mathbf{i}_{X_H} \omega := \omega(X_H, \cdot) = dH(\cdot) .$$

So the condition for a momentum map is that for all ξ

$$(332) \quad \omega(\xi_M, \cdot) = d(J(\xi))(\cdot) .$$

In Hamiltonian mechanics, it is common to write the $2n$ local coordinates q, p as ξ , i.e. to write

$$(333) \quad \xi^\alpha := q^\alpha, \quad \alpha = 1, \dots, n \quad ; \quad \xi^\alpha := p_{\alpha-n}, \quad \alpha = n+1, \dots, 2n .$$

So in order to express eq. 332 in local coordinates, let us temporarily write η for the arbitrary element of \mathfrak{g} . Then writing $\eta_M = (\eta_M)^\alpha \frac{\partial}{\partial \xi^\alpha}$ and $\omega_{\alpha\beta} := \omega(\frac{\partial}{\partial \xi^\alpha}, \frac{\partial}{\partial \xi^\beta})$, eq. 332 becomes

$$(334) \quad \omega_{\alpha\beta}(\eta_M)^\alpha = \frac{\partial J(\eta)}{\partial \xi^\beta} .$$

(3): *Components: an example:*

As discussed after eq. 313, we think of the collection of functions $J(\xi)$, as ξ varies through \mathfrak{g} , as the *components* of \mathbf{J} .

To take our standard example: the angular momentum of a particle in Euclidean space, in a state $x = (\mathbf{q}, \mathbf{p})$ is $\mathbf{J}(x) := \mathbf{q} \wedge \mathbf{p}$. Identifying $\mathfrak{so}(3)^*$ with \mathbb{R}^3 so that the natural pairing is given by the dot product (cf. (3) at the end of Section 4.5.2), we get that the component of $\mathbf{J}(x)$ around the axis $\xi \in \mathbb{R}^3$ is $\langle \mathbf{J}(x), \xi \rangle = \xi \cdot (\mathbf{q} \wedge \mathbf{p})$. The Hamiltonian vector field determined by this Hamiltonian function $x = (\mathbf{q}, \mathbf{p}) \mapsto \xi \cdot (\mathbf{q} \wedge \mathbf{p})$ is of course the infinitesimal generator of rotations about the ξ -axis. In Section 6.3, we will see more examples of momentum maps.

6.2 Conservation of momentum maps: Noether’s theorem

In ordinary Hamiltonian mechanics, we saw that Noether’s theorem had a simple expression as a “one-liner” based on the antisymmetry of the Poisson bracket: namely, in eq. 19, which was that for any scalar functions F, H

$$(335) \quad X_F(H) = \{H, F\} = 0 \quad \text{iff} \quad 0 = \{F, H\} = X_H(F) .$$

In words: the Hamiltonian H is constant under the flow induced by F iff F is a constant of the motion under the dynamical flow X_H .

More precisely, Section 2.1.3 vindicated this one-liner as expressing Noether’s theorem. For the one-liner respected the requirement that a symmetry should preserve the symplectic form (equivalently, the Poisson bracket), and not just (as in the left hand side of eq. 335) the Hamiltonian function H ; for, by Cartan’s magic formula, a vector field’s preserving the symplectic form was equivalent to its being locally Hamiltonian.

For Poisson manifolds, the equivalence corresponding to this last statement fails. That is, as we noted in (2) of Section 6.1.1: an infinitesimal Poisson automorphism need not be locally Hamiltonian.

Nevertheless, most of the “one-liner” approach to Noether’s theorem carries over to the framework of Poisson manifolds. In effect, we just restrict discussion to cases where the relevant Hamiltonian vector fields exist: recall our saying after (2) of Section 6.1.1 that we would concentrate on cases where all the ξ_M are globally Hamiltonian.

Thus, it is straightforward to show that for a Poisson manifold M , just as for symplectic manifolds: if $F, H \in \mathcal{F}(M)$, H is constant along the integral curves of X_F iff $\{H, F\} = 0$ iff F is constant along the integral curves of X_H . (We could

have proved this already in Section 5.2.2; but postponed it till now, when it will be used.)

With this result as a lemma, one immediately gets

Noether's theorem for Poisson manifolds Suppose that G acts canonically on a Poisson manifold M and has a momentum map $\mathbf{J} : M \rightarrow \mathfrak{g}^*$; and that H is invariant under ξ_M for all $\xi \in \mathfrak{g}$, i.e. $\{H, J(\xi)\} = \xi_M(H) = 0, \forall \xi \in \mathfrak{g}$; (cf. eq. 316). Then \mathbf{J} is a constant of the motion determined by H . That is:

$$(336) \quad \mathbf{J} \circ \phi_\tau = \mathbf{J}$$

where ϕ_τ is the flow of X_H .

Proof. By the lemma, the fact that $\{H, J(\xi)\} = \xi_M(H) = 0$ implies that $J(\xi)$ is constant along the flow of X_H . So by the definition of momentum map, eq. 323, the corresponding \mathfrak{g}^* -valued map \mathbf{J} is also a constant of the motion. ■

It follows immediately that H itself, and any distinguished function, is a constant of the motion. Besides, as remarked in (2) at the end of Section 6.1.1: a constant of the motion $J(\xi)$ is determined only up to an arbitrary choice of a distinguished function. Indeed, though this Chapter has set aside (ever since (iii) of Section 1.2) time-dependent functions: if one considers them, then there is here an arbitrary choice of a time-dependent distinguished function.

6.3 Examples

We give two familiar examples; and then, as a glimpse of the general power of the theory, two abstract examples (which will *not* be needed later on).

(1): *Total linear momentum of N particles:*

In (3) at the end of Section 4.1.0.1, we showed that the *left* cotangent lift of the action of the translation group \mathbb{R}^3 on $Q = \mathbb{R}^{3N}$ to $M = T^*\mathbb{R}^{3N}$, i.e. the left action corresponding to eq. 126, is

$$(337) \quad \Psi_{\mathbf{x}}(\mathbf{q}_i, \mathbf{p}^i) := T^*(\Phi_{-\mathbf{x}})(\mathbf{q}_i, \mathbf{p}^i) = (\mathbf{q}_i + \mathbf{x}, \mathbf{p}^i), \quad i = 1, \dots, N.$$

(Here we combine the discussions of examples (vi) and (ix) in Section 4.1.0.1)

To find the momentum map, we: (a) compute the infinitesimal generator ξ_M for an arbitrary element ξ of $\mathfrak{g} = \mathbb{R}^3$; and then (b) solve eq. 332, or in coordinates eq. 334.

(a) We differentiate eq. 337 with respect to \mathbf{x} in the direction ξ , getting

$$(338) \quad \xi_M(\mathbf{q}_i, \mathbf{p}^i) = (\xi, \dots, \xi, \mathbf{0}, \dots, \mathbf{0}).$$

(b) Any function $J(\xi)$ has Hamiltonian vector field

$$(339) \quad X_{J(\xi)}(\mathbf{q}_i, \mathbf{p}^i) = \left(\frac{\partial J(\xi)}{\partial \mathbf{p}^i}, -\frac{\partial J(\xi)}{\partial \mathbf{q}_i} \right);$$

so that the desired $J(\xi)$ with $X_{J(\xi)} = \xi_M$ solves

$$(340) \quad \frac{\partial J(\xi)}{\partial \mathbf{p}^i} = \xi \quad \text{and} \quad \frac{\partial J(\xi)}{\partial \mathbf{q}_i} = \mathbf{0}, \quad 1 \leq i \leq N.$$

Choosing constants so that J is linear, the solution is

$$(341) \quad J(\xi)(\mathbf{q}_i, \mathbf{p}^i) = (\Sigma_{i=1}^N \mathbf{p}^i) \cdot \xi, \quad \text{i.e.} \quad \mathbf{J}(\mathbf{q}_i, \mathbf{p}^i) = \Sigma_{i=1}^N \mathbf{p}^i;$$

i.e. the familiar total linear momentum.

(2): *Angular momentum of a single particle:*

$SO(3)$ acts on $Q = \mathbb{R}^3$ by $\Phi_A(\mathbf{q}) = A\mathbf{q}$. So the tangent (derivative) map is

$$(342) \quad T_{\mathbf{q}}\Phi_A : (\mathbf{q}, \mathbf{v}) \in T\mathbb{R}_{\mathbf{q}}^3 \mapsto (A\mathbf{q}, A\mathbf{v}) \in T\mathbb{R}_{A\mathbf{q}}^3.$$

As we saw in example (vii) of Section 4.1.0.1, the *left* cotangent lift of the action to $M = T^*\mathbb{R}^3$ (the lifted action “with g^{-1} ”, corresponding to eq. 126) is:

$$(343) \quad T_{A\mathbf{q}}^*(\Phi_{A^{-1}})(\mathbf{q}, \mathbf{p}) = (A\mathbf{q}, A\mathbf{p}).$$

To find the momentum map, we proceed in two stages, (a) and (b), as in example (1).

- (a) We differentiate eq. 343 with respect to A in the direction $\xi = \Theta(\omega) \in \mathfrak{so}(3)$, where $\omega \in \mathbb{R}^3$ and Θ is as in eq. 48 and 51. We get

$$(344) \quad \xi_M(\mathbf{q}, \mathbf{p}) = (\xi\mathbf{q}, \xi\mathbf{p}) = (\omega \wedge \mathbf{q}, \omega \wedge \mathbf{p}).$$

- (b) So the desired $J(\xi)$ is the solution linear in ξ to the Hamilton’s equations

$$(345) \quad \frac{\partial J(\xi)}{\partial \mathbf{p}} = \xi\mathbf{q} \quad \text{and} \quad \frac{\partial J(\xi)}{\partial \mathbf{q}} = -\xi\mathbf{p}.$$

So a solution is given by

$$(346) \quad J(\xi)(\mathbf{q}, \mathbf{p}) = (\xi\mathbf{q}) \cdot \mathbf{p} = (\omega \wedge \mathbf{q}) \cdot \mathbf{p} = (\mathbf{q} \wedge \mathbf{p}) \cdot \omega,$$

so that

$$(347) \quad \mathbf{J}(\mathbf{q}, \mathbf{p}) = \mathbf{q} \wedge \mathbf{p},$$

i.e. the familiar angular momentum.

(3): *Dual of a Lie algebra homomorphism:*

We begin by stating a *Lemma*, which we will not prove; for details cf. Marsden and Ratiu [1999, 10.7.2, 372]. Namely: let G, H be Lie groups and let $\alpha : \mathfrak{g} \rightarrow \mathfrak{h}$ be a linear map between their Lie algebras. Then α is a Lie algebra homomorphism iff its dual $\alpha^* : \mathfrak{h}^* \rightarrow \mathfrak{g}^*$ is a (linear) Poisson map (where $\mathfrak{h}^*, \mathfrak{g}^*$ are equipped with their natural Lie-Poisson brackets as in Section 5.2.4).

Now let G, H be Lie groups, let $A : H \rightarrow G$ be a Lie group homomorphism, and let $\alpha : \mathfrak{h} \rightarrow \mathfrak{g}$ be the induced Lie algebra homomorphism; so that by the Lemma, $\alpha^* : \mathfrak{g}^* \rightarrow \mathfrak{h}^*$ is a Poisson map. We will prove that α^* is also a momentum map for the action of H on \mathfrak{g}^* given by, with $h \in H, x \in \mathfrak{g}^*$:

$$(348) \quad \Phi(h, x) \equiv h \cdot x := Ad_{A(h)}^* x .$$

Proof. We first recall the adjoint and co-adjoint actions $Ad_g : \mathfrak{g} \rightarrow \mathfrak{g}$ and $Ad_g^* : \mathfrak{g}^* \rightarrow \mathfrak{g}^*$; in particular, eq. 191. So the action in eq. 348 is:

$$(349) \quad \forall x \in \mathfrak{g}^*, \forall \xi \in \mathfrak{g} : \langle h \cdot x; \xi \rangle = \langle x; Ad_{A(h)} \xi \rangle .$$

As usual, we compute for $\eta \in \mathfrak{h}$, the infinitesimal generator $\eta_{\mathfrak{g}^*}$ at $x \in \mathfrak{g}^*$ by differentiating eq. 349 with respect to h at e in the direction $\eta \in \mathfrak{h}$. We get (cf. eq. 198):

$$(350) \quad \langle \eta_{\mathfrak{g}^*}(x); \xi \rangle = - \langle x; ad_{\alpha(\eta)} \xi \rangle = \langle ad_{\alpha(\eta)}^*(x); \xi \rangle .$$

We define $\mathbf{J}(x) := \alpha^*(x)$: that is,

$$(351) \quad J(\eta)(x) \equiv \langle \mathbf{J}(x); \eta \rangle := \langle \alpha^*(x); \eta \rangle \equiv \langle x; \alpha(\eta) \rangle ;$$

which implies

$$(352) \quad \nabla_x J(\eta) = \alpha(\eta) .$$

Now we recall that Hamilton's equations for $J(\eta)$ as the Hamiltonian are (cf. eq. 261)

$$(353) \quad \dot{x} \equiv X_{J(\eta)}(x) = ad_{\nabla_x J(\eta)}^*(x) .$$

Combining eq. 350 to eq. 353, we get:

$$(354) \quad X_{J(\eta)}(x) = ad_{\alpha(\eta)}^*(x) = \eta_{\mathfrak{g}^*}(x) ;$$

proving that $\mathbf{J}(x) := \alpha^*(x)$ is a momentum map. ■

(4): *Momentum maps for subgroups:*

Assume that $\mathbf{J} : M \rightarrow \mathfrak{g}^*$ is a momentum map for a canonical left action of G on M ; and let $H < G$ be a subgroup of G . Then H also acts canonically on M , and this action has as a momentum map the restriction of \mathbf{J} 's values to $\mathfrak{h} \subset \mathfrak{g}$. That is: the map

$$(355) \quad \mathbf{J}_H : M \rightarrow \mathfrak{h}^* \quad \text{given by} \quad \mathbf{J}_H(x) := \mathbf{J}(x) |_{\mathfrak{h}} .$$

For the canonical action of G ensures that if $\eta \in \mathfrak{h} \subset \mathfrak{g}$, then $\eta_M = X_{J(\eta)}$. Then $J_H(\eta) := J(\eta) \forall \eta \in \mathfrak{h}$ defines a momentum map for H 's action. That is

$$(356) \quad \forall x \in M, \forall \eta \in \mathfrak{h} : \langle \mathbf{J}_H(x); \eta \rangle = \langle \mathbf{J}(x); \eta \rangle .$$

6.4 Equivariance of momentum maps

In (1) of Section 4.2, we defined the general notion of an equivariant map $f : M \rightarrow N$ between manifolds as one that respects the actions of a group G on M and on N : eq. 144. We now develop an especially important case of this notion: the equivariance of momentum maps $\mathbf{J} : M \rightarrow \mathfrak{g}^*$, where the action on \mathfrak{g}^* is the co-adjoint action, eq. 192.

For us, this notion will have two main significances:

- (i) many momentum maps that occur in examples are equivariant in this sense;
- (ii) equivariance has various theoretical consequences: in particular, momentum maps for cotangent lifted actions are always equivariant (Section 6.5), and equivariance is crucial in theorems about reduction (Section 7).

In this Section, we will glimpse these points by:

- (i) defining the notion, and remarking on a weakened differential version of the notion (Section 6.4.1);
- (ii) proving that equivariant momentum maps are Poisson (Section 6.4.2).

6.4.1 Equivariance and infinitesimal equivariance

Let Φ be a canonical left action of G on M , and let $\mathbf{J} : M \rightarrow \mathfrak{g}^*$ be a momentum map for it. We say \mathbf{J} is *equivariant* if for all $g \in G$

$$(357) \quad \mathbf{J} \circ \Phi_g = Ad_{g^{-1}}^* \circ \mathbf{J} ;$$

cf. eq. 144 and the definition of co-adjoint action, eq. 193:

$$(358) \quad \begin{array}{ccc} M & \xrightarrow{\mathbf{J}} & \mathfrak{g}^* \\ \uparrow \Phi_g & & \uparrow Ad_{g^{-1}}^* \\ M & \xrightarrow{\mathbf{J}} & \mathfrak{g}^* \end{array}$$

An equivalent formulation arises by considering that we can add to the commutative square in eq. 358 the two commutative triangles:

$$(359) \quad M \xrightarrow{J(\xi)} \mathbb{R} \quad \text{is} \quad M \xrightarrow{\mathbf{J}} \mathfrak{g}^* \xrightarrow{|\xi} \mathbb{R} ;$$

representing the fact that $J(\xi)(x) = \mathbf{J}(x)(\xi)$; and

$$(360) \quad \mathfrak{g}^* \xrightarrow{|\xi} \mathbb{R} \quad \text{is} \quad \mathfrak{g}^* \xrightarrow{Ad_{g^{-1}}^*} \mathfrak{g}^* \xrightarrow{|\xi} \mathbb{R} ;$$

representing the fact that for all $\eta \in \mathfrak{g}^*$

$$(361) \quad \langle Ad_{g^{-1}}^*(\eta); Ad_g(\xi) \rangle = \langle \eta; Ad_{g^{-1}} Ad_g(\xi) \rangle \equiv \langle \eta; \xi \rangle .$$

Eq.s 359 and 360 imply that an equivalent formulation of equivariance is that for all $x \in M, g \in G$ and $\xi \in \mathfrak{g}$ (and with $g \cdot x \equiv \Phi_g(x)$)

$$(362) \quad \mathbf{J}(g \cdot x)(Ad_g \xi) \equiv J(Ad_g \xi)(g \cdot x) = J(\xi)(x) \equiv \mathbf{J}(x)(\xi) .$$

In (2) of Section 4.4, we differentiated the general notion of an equivariant map, and got the weaker differential notion that the infinitesimal generators ξ_M and ξ_N of the actions of G on M and on N are f -related.

Here also we can differentiate equivariance, and get the notion of *infinitesimal equivariance*. But I will not go into details since:

- (i) we will not need the notion, not least because (as mentioned above), many momentum maps are equivariant;
- (ii) under certain common conditions (e.g. the group G is compact, or is connected) an infinitesimally equivariant momentum map can always be replaced by an equivariant one.

So let it suffice to say that infinitesimal equivariance *is* theoretically important. In particular, the result eq. 317, viz.

$$(363) \quad X_{J([\xi, \eta])} = X_{\{J(\xi), J(\eta)\}_M}$$

implies that

$$(364) \quad \Sigma(\xi, \eta) := J([\xi, \eta]) - \{J(\xi), J(\eta)\}_M$$

is a distinguished function on the Poisson manifold M , and so constant on every symplectic leaf.

This makes it natural to ask when $\Sigma \equiv 0$. After all, cf. eq. 322. Both $\xi \mapsto \xi_M$ and $F \mapsto X_F$ are Lie algebra anti-homomorphisms. So it is natural to ask whether J is a Lie algebra homomorphism, i.e. whether $\Sigma = 0$. And it turns out that infinitesimal equivariance is equivalent to $\Sigma = 0$.

6.4.2 Equivariant momentum maps are Poisson

The following result is important, both as a general method of finding canonical maps between Poisson manifolds, and for the Lie-Poisson reduction theorem of Section 7.

Equivariant momentum maps are Poisson Let $\mathbf{J} : \mathbf{M} \rightarrow \mathfrak{g}^*$ be an equivariant momentum map for a canonical left action of G on a Poisson manifold M . Then \mathbf{J} is a Poisson map: for all $F_1, F_2 \in \mathcal{F}(\mathfrak{g}^*)$,

$$(365) \quad \mathbf{J}^* \{F_1, F_2\}_{\mathfrak{g}^*} = \{\mathbf{J}^* F_1, \mathbf{J}^* F_2\}_M \quad ; \text{ i.e. } \quad \{F_1, F_2\}_{\mathfrak{g}^*} \circ \mathbf{J} = \{\mathbf{J} \circ F_1, \mathbf{J} \circ F_2\}_M .$$

Proof. We will relate (i) the left hand side, then (ii) the right hand side of eq. 365 to J ; and finally we will use the fact that the Poisson bracket on M depends only on the values of the first derivatives.

- (i) Let $x \in M, \alpha = \mathbf{J}(x) \in \mathfrak{g}^*$; and let $\xi = \nabla F_1$ and $\eta = \nabla F_2$ evaluated at α , so that $\xi, \eta \in \mathfrak{g}^{**} = \mathfrak{g}$. Then

$$(366) \quad \{F_1, F_2\}_{\mathfrak{g}^*}(\mathbf{J}(x)) \equiv \langle \alpha; [\nabla F_1, \nabla F_2] \rangle = \langle \alpha; [\xi, \eta] \rangle = J([\xi, \eta])(x) = \{J(\xi), J(\eta)\}(x);$$

where the third equation just applies the definition of \mathbf{J} , eq. 323, and the fourth equation uses (infinitesimal) equivariance.

- (ii) We show that $(F_1 \circ \mathbf{J})(x)$ and $J(\xi)(x)$ have equal x -derivatives. For any $x \in M$ and $v_x \in T_x M$

$$(367) \quad \mathbf{d}(F_1 \circ \mathbf{J})(x) \cdot v_x = \mathbf{d}F_1(\alpha) \cdot T_x \mathbf{J}(v_x) = \langle T_x \mathbf{J}(v_x); \nabla F_1 \rangle = \mathbf{d}J(\xi)(x) \cdot v_x;$$

where the first equation uses the chain rule, and the last uses the definition of \mathbf{J} , eq. 323 and the fact that $\xi = \nabla F_1$.

Finally, since the Poisson bracket on M depends only on the values of the first derivatives, we infer from eq. 367 that

$$(368) \quad \{F_1 \circ \mathbf{J}, F_2 \circ \mathbf{J}\}(x) = \{J(\xi), J(\eta)\}(x).$$

Combining this with (i), the result follows. ■

6.5 Momentum maps on cotangent bundles

Let a Lie group G act on a manifold (“configuration space”) Q . We saw in Section 4.1.0.1 that this action can be *lifted* to the cotangent bundle T^*Q ; cf. eq.s 121, 124 and 126. In this Section, we focus on momentum maps for such *cotangent lift* actions. We shall see that any such action has an equivariant momentum map, for which there is an explicit general formula. The general theory (Sections 6.5.1, 6.5.2) will need just one main new notion, the *momentum function*. We end with some examples (Section 6.5.3).

6.5.1 Momentum functions

Given a manifold Q and its vector fields $\mathcal{X}(Q)$, we define the map

$$(369) \quad \mathcal{P} : \mathcal{X}(Q) \rightarrow \mathcal{F}(T^*Q) \quad \text{by} : \quad (\mathcal{P}(X))(\alpha_q) := \langle \alpha_q; X(q) \rangle$$

for $q \in Q, X \in \mathcal{X}(Q)$ and $\alpha_q \in T_q^*Q$. Here, α_q is, strictly speaking, a point in the cotangent bundle above the base-point $q \in Q$: so α_q can be written as (q, α) with α a covector at q , i.e. $\alpha \in T_q^*Q$. But as we mentioned just before defining cotangent lifts (eq. 121): it is harmless to (follow many presentations and) conflate a point in T^*Q , i.e. a pair $(q, \alpha), q \in Q, \alpha \in T_q^*Q$, with its form α , provided we keep track of the q by writing the form as α_q .

$\mathcal{P}(X)$, as defined by eq. 369, is called the *momentum function* of X . In coordinates, $\mathcal{P}(X)$ is given by

$$(370) \quad \mathcal{P}(X)(q^i, p_i) = X^j(q^i)p_j$$

where we sum on $j = 1, \dots, n := \dim Q$. (So NB: This \mathcal{P} is different from that in $\mathcal{P}(M)$, the infinitesimal Poisson automorphisms of M , discussed in Section 6.1.1.)

We also denote by $\mathcal{L}(T^*Q)$ the space of smooth functions $F : T^*Q \rightarrow \mathbb{R}$ that are linear on fibres of T^*Q : i.e. writing the bundle points $\alpha_q, \beta_q \in T_q^*Q$ as (q, α) and (q, β) , we have for $\lambda, \mu \in \mathbb{R}$

$$(371) \quad F(q, (\lambda\alpha + \mu\beta)) = \lambda F((q, \alpha)) + \mu F((q, \beta)).$$

So functions F, H that are in $\mathcal{L}(T^*Q)$ can be written in coordinates as (summing on $i = 1, \dots, n$)

$$(372) \quad F(q, p) = X^i(q)p_i \quad \text{and} \quad H(q, p) = Y^i(q)p_i$$

for functions X^i and Y^i ; and so any momentum function $\mathcal{P}(X)$ is in $\mathcal{L}(T^*Q)$.

One readily checks that the standard Poisson bracket (from T^*Q 's symplectic structure, Section 2.1.1) of such an F and H is also linear on the fibres of T^*Q . In fact, eq. 372 implies

$$(373) \quad \{F, H\}(q, p) := \frac{\partial F}{\partial q^j} \frac{\partial H}{\partial p_j} - \frac{\partial H}{\partial q^j} \frac{\partial F}{\partial p_j} = \left(\frac{\partial X^i}{\partial q^j} Y^j - \frac{\partial Y^i}{\partial q^j} X^j \right).$$

So $\mathcal{L}(T^*Q)$ is a Lie subalgebra of $\mathcal{F}(T^*Q)$.

The next result summarizes how momentum functions relate $\mathcal{X}(Q)$ and Hamiltonian vector fields on T^*Q to $\mathcal{L}(T^*Q)$.

Three (anti)-isomorphic Lie algebras The two Lie algebras

- (i) $(\mathcal{X}(Q), [,])$ of vector fields on Q ;
- (ii) Hamiltonian vector fields X_F on T^*Q with $F \in \mathcal{L}(T^*Q)$

are isomorphic. And each is anti-isomorphic to

- (iii) $(\mathcal{L}(T^*Q), \{, \})$.

In particular, the map \mathcal{P} is an anti-isomorphism from (i) to (iii), so that we have

$$(374) \quad \{\mathcal{P}(X), \mathcal{P}(Y)\}_{T^*Q} = -\mathcal{P}([X, Y]).$$

Proof. Since $\mathcal{P}(X) : T^*Q \rightarrow \mathbb{R}$ is linear on fibres, \mathcal{P} maps $\mathcal{X}(Q)$ into $\mathcal{L}(T^*Q)$. \mathcal{P} is also onto $\mathcal{L}(T^*Q)$: given $F \in \mathcal{L}(T^*Q)$, we can define $X(F) \in \mathcal{X}(Q)$ by

$$(375) \quad \langle \alpha_q; X(F)(q) \rangle := F(\alpha_q) \quad \forall \alpha_q \in T_q^*Q$$

so that $\mathcal{P}(X(F)) = F$. \mathcal{P} is linear and $\mathcal{P}(X) = 0$ implies that $X = 0$. Also, eq. 374 follows immediately by comparing eq. 373 with the Lie bracket of $X, Y \in \mathcal{X}(Q)$; cf. eq. 55. So \mathcal{P} is an anti-isomorphism from $(\mathcal{X}Q, [,])$ to $(\mathcal{L}(T^*Q), \{, \})$.

The map

$$(376) \quad F \in (\mathcal{L}(T^*Q), \{, \}) \mapsto X_F \in (\{X_F \mid F \in \mathcal{L}(T^*Q)\}, [,])$$

is surjective by definition. It is a Lie algebra anti-homomorphism, by eq. 60 (i.e. result (2) in Section 3.2.2). And if $X_F = 0$, then F is constant on T^*Q ; and hence $F \equiv 0$ since F is linear on the fibres (cf. eq. 371). ■

6.5.2 Momentum maps for cotangent lifted actions

We begin this Subsection with a result relating the Hamiltonian flow on T^*Q induced by the momentum function $\mathcal{P}(X)$ to the Hamiltonian flow on X induced by X . From this result, our main result — the guarantee of an equivariant momentum map for a cotangent lifted action, and an explicit formula for it — will follow directly.

The Hamiltonian flow of a momentum function Let $X \in \mathcal{X}(Q)$ have flow ϕ_τ on Q ; cf. Section 3.1.2.2. Then the flow of $X_{\mathcal{P}(X)}$ on T^*Q is $T^*\phi_{-\tau}$. That is: the flow of $X_{\mathcal{P}(X)}$ is the cotangent lift (Section 4.1.0.1) of $\phi_{-\tau}$, as given by the diagram, with π_Q the canonical projection:

$$(377) \quad \begin{array}{ccc} Q & \xrightarrow{\phi_\tau} & Q \\ \uparrow \pi_Q & & \uparrow \pi_Q \\ T^*Q & \xrightarrow{T^*\phi_{-\tau}} & T^*Q \end{array}$$

Proof. We differentiate the relation in eq. 377, i.e.

$$(378) \quad \pi_Q \circ T^*\phi_{-\tau} = \phi_\tau \circ \pi_Q$$

at $\tau = 0$ to get

$$(379) \quad T\pi_Q \circ Y = X \circ \pi_Q \quad \text{with} \quad \forall \alpha_q \in T_q^*Q, \quad Y(\alpha_q) = \left. \frac{d}{d\tau} \right|_{\tau=0} T^*\phi_{-\tau}(\alpha_q);$$

i.e. $T^*\phi_{-\tau}$ is the flow of Y .

Now we will show that $Y = X_{\mathcal{P}(X)}$, using eq. 379 and the geometrical formulation of Hamiltonian mechanics of Section 2.1, especially Cartan’s magic formula, eq. 20, applied to the canonical one-form $\theta \equiv \theta_H$ (defined by eq. 8 and 9).

We reported (at the start of (2) of Section 4.1.0.1) that the cotangent lift $T^*\phi_{-\tau}$ preserves $\theta \equiv \theta_H$ on T^*Q . So $\mathcal{L}_Y\theta = 0$. Then the definition of ω as the negative exterior derivative of θ , and Cartan’s magic formula, eq. 20, yields

$$(380) \quad \mathbf{i}_Y\omega = -\mathbf{i}_Y\mathbf{d}\theta = \mathbf{d}\mathbf{i}_Y\theta.$$

On the other hand, we also have

$$(381) \quad \mathbf{i}_Y\theta(\alpha_q) \equiv \langle \theta(\alpha_q); Y(\alpha_q) \rangle = \langle \alpha_q; T\pi_Q(Y(\alpha_q)) \rangle = \langle \alpha_q; X(q) \rangle = \mathcal{P}(X)(\alpha_q)$$

where the second equation applies the definition of the canonical one-form (eq. 8), the third applies eq. 379, and the fourth applies the definition eq. 369 of momentum functions.

Combining eq. 380 and 381, we have:

$$(382) \quad \mathbf{i}_Y \omega = \mathbf{d}\mathcal{P}(X)$$

which is Hamilton's equations (eq. 15) telling us that $Y = X_{\mathcal{P}(X)}$. ■

Accordingly the Hamiltonian vector field $X_{\mathcal{P}(X)}$ on T^*Q is called the *cotangent lift* of $X \in \mathcal{Q}$ to T^*Q . In local coordinates, we can write, by combining eq. 15 and 370

$$(383) \quad X_{\mathcal{P}(X)} = \frac{\partial \mathcal{P}(X)}{\partial p_i} \frac{\partial}{\partial q^i} - \frac{\partial \mathcal{P}(X)}{\partial q^i} \frac{\partial}{\partial p_i} = X^i \frac{\partial}{\partial q^i} - \frac{\partial X^i}{\partial q^j} p^i \frac{\partial}{\partial p_j} .$$

Note in particular that, combining the usual sign-change between Lie algebras and Poisson brackets (eq. 60) with the sign-change for momentum functions (eq. 374), we have

$$(384) \quad [X_{\mathcal{P}(X)}, X_{\mathcal{P}(Y)}] = -X_{\{\mathcal{P}(X), \mathcal{P}(Y)\}} = -X_{-\mathcal{P}([X, Y])} = X_{\mathcal{P}([X, Y])} .$$

We can now readily prove our main result guaranteeing, and giving a formula for, equivariant momentum maps.

Equivariant momentum maps Let G act on the left on Q and so by cotangent lift on T^*Q . The cotangent lifted action has an equivariant momentum map $\mathbf{J} : T^*Q \rightarrow \mathfrak{g}^*$ given by

$$(385) \quad \langle \mathbf{J}(\alpha_q); \xi \rangle = \langle \alpha_q; \xi_Q(q) \rangle \equiv \mathcal{P}(\xi_Q)(\alpha_q) .$$

In coordinates q^i, p_i on T^*Q and ξ^a on \mathfrak{g} , and with $\xi_Q^i = \xi^a A_a^i$ the components of ξ_Q , this reads

$$(386) \quad J_a \xi^a = p_i \xi_Q^i = p_i A_a^i \xi^a$$

so that $J_a(q, p) = p_i A_a^i(q)$.

Proof. The preceding result tells us that for any $\xi \in \mathfrak{g}$, the infinitesimal generator of the cotangent lifted action on T^*Q is $\xi_{T^*Q} \equiv X_{\mathcal{P}(\xi_Q)}$. So a momentum map for this action is given by

$$(387) \quad J(\xi) = \mathcal{P}(\xi_Q) .$$

This gives eq. 385, just by applying the definitions of the momentum map \mathbf{J} (eq. 323) and of momentum function (eq. 369).

To prove equivariance, we argue as follows:

$$(388) \quad \langle \mathbf{J}(g \cdot \alpha_q); \xi \rangle = \langle (g \cdot \alpha_q); \xi_Q(g \cdot q) \rangle$$

$$(389) = \langle \alpha_q; (T\Phi_{g^{-1}})\xi_Q(g \cdot q) \rangle \equiv \langle \alpha_q; (T_{g \cdot q}\Phi_{g^{-1}} \circ \xi_Q \circ \Phi_g)(q) \rangle$$

$$(390) = \langle \alpha_q; (\Phi_g^* \xi_Q)(q) \rangle$$

$$(391) = \langle \alpha_q; (Ad_{g^{-1}} \xi)_Q(q) \rangle$$

$$(392) = \langle \mathbf{J}(\alpha_q); Ad_{g^{-1}} \xi \rangle = \langle Ad_{g^{-1}}^* (\mathbf{J}(\alpha_q)); \xi \rangle .$$

Here we have applied in succession: (i) eq. 385; (ii) the fact that $g \cdot \alpha_q$ is short for $T^*(\Phi_{g^{-1}})(\alpha_q)$, cf. eq. 126 and 121; (iii) the definition of pullback, cf. eq. 172; (iv) result [2], eq. 167, of Section 4.5.1; (v) eq. 385 again; and finally, (vi) the fact that Ad^* is the adjoint of Ad , cf. eq. 191. ■

6.5.3 Examples

We discuss first our familiar examples, linear and angular momentum i.e. (1) and (2) from Section 6.3; and then the cotangent lift of left and right translations on G — an example motivated by Section 4.6’s description of kinematics on a Lie group G .

(1): *Total linear momentum of N particles:*

Since the translation group \mathbb{R}^3 acts on $Q := \mathbb{R}^{3N}$ by $\Phi(\mathbf{x}, (\mathbf{q}_i)) = (\mathbf{q}_i + \mathbf{x})$, the infinitesimal generator on Q is

$$(393) \quad \xi_{\mathbb{R}^{3N}}(\mathbf{q}_i) = (\xi, \dots, \xi) \quad (\xi \text{ } N \text{ times})$$

Applying eq. 385, the equivariant momentum map is given by

$$(394) \quad J(\xi)(\mathbf{q}_i, \mathbf{p}^i) = (\Sigma_{i=1}^N \mathbf{p}^i) \cdot \xi, \quad \text{i.e.} \quad \mathbf{J}(\mathbf{q}_i, \mathbf{p}^i) = \Sigma_{i=1}^N \mathbf{p}^i;$$

agreeing with our previous solution, eq. 341, based on the differential equation eq. 334.

(2): *Angular momentum of a single particle:*

$SO(3)$ acts on \mathbb{R}^3 by $\Phi(A, \mathbf{q}) = A\mathbf{q}$. Writing $\xi \in \mathfrak{so}(3)$ as $\xi = \Theta\omega$ (cf. eq. 47, 51 and 105), the infinitesimal generator is

$$(395) \quad \xi_{\mathbb{R}^3}(\mathbf{q}) = \xi\mathbf{q} = \omega \wedge \mathbf{q}.$$

So applying eq. 385, the equivariant momentum map $\mathbf{J} : T^*\mathbb{R}^3 \rightarrow \mathfrak{so}(3) \cong \mathbb{R}^3$ is given by

$$(396) \quad \langle \mathbf{J}(\mathbf{q}, \mathbf{p}); \omega \rangle = \langle \mathbf{p}; \omega \wedge \mathbf{q} \rangle = \mathbf{p} \cdot (\omega \wedge \mathbf{q}) = \omega \cdot (\mathbf{q} \wedge \mathbf{p}), \quad \text{i.e.} \quad \mathbf{J}(\mathbf{q}, \mathbf{p}) = \mathbf{q} \wedge \mathbf{p};$$

agreeing with our previous solution, eq. 347, based on the differential equation eq. 334.

(3): *The cotangent lift of left and right translations on G :*

Recalling eq. 157, viz. that the infinitesimal generator of left translation is

$$(397) \quad \xi_G(g) = (T_e R_g)\xi,$$

a *right*-invariant vector field, and applying eq. 385, we see that the momentum map $\mathbf{J}_L : T^*G \rightarrow \mathfrak{g}^*$ for the cotangent lift of left translation is given by

$$(398) \quad \langle \mathbf{J}_L(\alpha_g); \xi \rangle = \langle \alpha_g; \xi_G(g) \rangle = \langle \alpha_g; (T_e R_g)\xi \rangle = \langle (T_e^* R_g)(\alpha_g); \xi \rangle$$

where the last equation applies the definition of the cotangent lift eq. 121. That is: the equivariant momentum map is

$$(399) \quad \mathbf{J}_L(\alpha_g) = T_e^* R_g(\alpha_g) .$$

In words: the momentum map \mathbf{J}_L of the cotangent lift of left translation is the cotangent lift of right translation.

In a similar way, we could consider right translation: $R_g : h \mapsto hg$. Right translation defines a right action on G , has $\xi_G(g) = (T_e L_g)\xi$ as its infinitesimal generator, and so has

$$(400) \quad \mathbf{J}_R : T^*G \rightarrow \mathfrak{g}^* ; \quad \mathbf{J}_R(\alpha_g) := T_e^* L_g(\alpha_g)$$

as the momentum map of its cotangent lift. Note that this momentum map is equivariant with respect to Ad_g^* : which, as discussed after eq. 191, is a *right* action.

7 REDUCTION

7.1 Preamble

In this final Section, the themes of Section 2 onwards come together — at last! As announced in Section 5.1, we will concentrate on proving what is nowadays called the *Lie-Poisson reduction theorem*: that is, the isomorphism of Poisson manifolds

$$(401) \quad T^*G/G \cong \mathfrak{g}^* .$$

Here the quotient of T^*G is by the cotangent lift of G 's action on itself by left translation.

As it happens, this Chapter's main sources (i.e. [Abraham and Marsden, 1978; Arnold, 1989; Olver, 2000; Marsden and Ratiu, 1999]) do not contain what is surely the most direct proof of this result. So we give it in Section 7.2. The result will follow directly from four previous main results, one from Section 5 and three from Section 6.

'Directly', but for one wrinkle! This relates to "flipping" between left and right translation, and their various lifts. In short: the four previous results show that T^*G/G is isomorphic as a Poisson manifold, *not* to \mathfrak{g}^* with the Lie-Poisson bracket familiar since eq. 255 and 258, but instead to \mathfrak{g}^* equipped with this bracket's *negative*, i.e. equipped with

$$(402) \quad \{F, H\}_-(x) := - \langle x; [\nabla F(x), \nabla H(x)] \rangle , \quad x \in \mathfrak{g}^* .$$

But we shall (mercifully!) not reproduce, with minus signs appropriately added, our entire discussion of the Lie-Poisson bracket that ensued after eq. 255; (exercise for the reader!).

To avoid ambiguity, we shall sometimes write \mathfrak{g}_+^* for \mathfrak{g}^* equipped with the positive Lie-Poisson bracket of eq. 258; and \mathfrak{g}_-^* for \mathfrak{g}^* equipped with the negative Lie-Poisson bracket of eq. 402.

In fact, it will be clearest from now on, to treat right actions on a par with left actions; despite our previous emphasis on the latter. This will mean that we will

also treat right-invariant vector fields (and another notion of right-invariance defined in Section 7.3.1) on a par with left-invariant vector fields (and Section 7.3.1's corresponding new notion of left-invariance). Indeed, we have already glimpsed this would be necessary in:

- (i) Section 4.4's result that the infinitesimal generator of left translation is a right-invariant vector field, and vice versa (eq. 157, 158); and its corollaries in Example (3) of Section 6.5.3, that
- (ii) the momentum map \mathbf{J}_L of the cotangent lift of left translation is the cotangent lift of right translation; (eq. 399); and
- (iii) the momentum map \mathbf{J}_R of the cotangent lift of right translation is the cotangent lift of left translation; (eq. 400).

So by the end of Section 7.2, we will have a short proof of the Lie-Poisson reduction theorem. But (as often happens), the most direct proof does not give very much information about the situation. So in Section 7.3 we give more information (following Marsden and Ratiu (1999)). Then in Section 7.4, we discuss the reduction of dynamics (as against Poisson structure) from T^*G to \mathfrak{g}^* .

Finally, in Section 7.5 we state another reduction theorem, which is cast in terms of symplectic, not Poisson, manifolds — but which uses several notions from Section 3, such as free and proper actions, and isotropy groups. But we do not prove this theorem: we include it mostly in order to emphasize our previous remark, that (despite its length!) this Chapter just scratches the surface of the subject. We also discuss the relation between it and the Lie-Poisson reduction theorem.

7.2 The Lie-Poisson Reduction Theorem

First we recall from the end of Section 4.6.2 (eq. 225) that $\bar{\lambda} : T^*G \rightarrow G \times \mathfrak{g}^*$ is an equivariant map relating the cotangent lifted left action of left translation on T^*G to the G -action on $G \times \mathfrak{g}^*$ given just by left translation on the first component. So we passed to the quotients, and defined $\hat{\lambda} : T^*G/G \rightarrow (G \times \mathfrak{g}^*)/G$ by eq. 228, viz.

$$(403) \quad \hat{\lambda} : \text{Orb}(\alpha) \equiv \{\beta \in T^*G \mid \beta = T^*L_{h^{-1}}(\alpha), \text{ some } h \in G\} \mapsto \\ \text{Orb}(\bar{\lambda}(\alpha)) \equiv \{(hg, (T_e^*L_g)(\alpha)) \mid \text{some } h \in G\} \equiv \\ \{(h, (T_e^*L_g)\alpha) \mid \text{some } h \in G\}.$$

where $\alpha \in T_g^*G$, so that $T^*L_{h^{-1}}\alpha \in T_{hg}^*G$. Finally, we identified $(G \times \mathfrak{g}^*)/G$ with \mathfrak{g}^* , so that the diffeomorphism $\hat{\lambda}$ maps T^*G/G to \mathfrak{g}^* , as in eq. 229:

$$(404) \quad \hat{\lambda} : \text{Orb}(\alpha) \equiv \{\beta \in T^*G \mid \beta = T^*L_{h^{-1}}(\alpha), \text{ some } h \in G\} \in T^*G/G \mapsto \\ (T_e^*L_g)(\alpha) \in \mathfrak{g}^*.$$

So now, we are to show that the diffeomorphism $\hat{\lambda} : T^*G/G \rightarrow \mathfrak{g}^*$ is a Poisson map, in the sense of eq. 283 (Section 5.3.2). So we need to show:

- (i) T^*G/G is a Poisson manifold;
- (ii) $\hat{\lambda}$ maps (i)'s Poisson structure on T^*G/G to that of \mathfrak{g}^* . In fact, as announced in Section 7.1, $\hat{\lambda}$ maps on to the Poisson structure of \mathfrak{g}_-^* , i.e. as given by eq. 402.

Prima facie, there could be a judicious choice to be made about (i), i.e. about how to define the Poisson structure on T^*G/G , so as to secure (ii), i.e. so that $\hat{\lambda}$ respects the Poisson structure. But in fact our previous work gives a pre-eminently obvious choice — which works. Namely: we use the Poisson structure induced on T^*G/G by the Poisson reduction theorem of Section 5.5. The result follows directly by combining with this theorem, three results from Section 6:

- (i) that equivariant momentum maps are Poisson; eq. 365 in Section 6.4.2;
- (ii) that a cotangent lifted left action has an equivariant momentum map; eq. 385 in Section 6.5.2;
- (iii) that the momentum maps of the cotangent lifts of left and right translation on G are $\mathbf{J}_L = T_e^*R_g$ and $\mathbf{J}_R = T_e^*L_g$; eq. 399 and 400 in Section 6.5.3.

In particular, combining (i)-(iii): one deduces (exercise!) that $\mathbf{J}_R = T_e^*L_g$ is equivariant with respect to Ad_g^* , and so Poisson with respect to the negative Lie-Poisson bracket (eq. 402's bracket) on \mathfrak{g}^* . That is: it is Poisson with the codomain \mathfrak{g}_-^* .

Thus we have the

Lie-Poisson reduction theorem The diffeomorphism $\hat{\lambda} : T^*G/G \rightarrow \mathfrak{g}_-^*$:

$$(405) \quad \hat{\lambda} : \text{Orb}(\alpha) \equiv \{\beta \in T^*G \mid \beta = T^*L_{h^{-1}}(\alpha), \text{ some } h \in G\} \in T^*G/G \mapsto (T_e^*L_g)(\alpha) \in \mathfrak{g}_-^*$$

is Poisson.

Proof. First, eq. 405 means we have a commutative triangle. For with $\pi : T^*G \rightarrow T^*G/G$ the canonical projection, the momentum map $\mathbf{J}_R : T^*G \rightarrow \mathfrak{g}_-^*, \alpha_g \mapsto (T_e^*L_g)\alpha_g$ is equal to $\hat{\lambda} \circ \pi$:

$$(406) \quad T^*G \xrightarrow{\pi} T^*G/G \xrightarrow{\hat{\lambda}} \mathfrak{g}_-^* .$$

Since left translation is a diffeomorphism of G , and the cotangent lift of any diffeomorphism of a manifold to its cotangent bundle is symplectic (cf. after eq. 120 in Section 4.1.0.1), the Poisson reduction theorem of Section 5.5 applies. That is, there is a unique Poisson structure on T^*G/G such that π is Poisson. We also know from eq. 385, 365 and 400 that $\mathbf{J}_R = T_e^*L_g$ is Poisson with respect to eq. 402's bracket on \mathfrak{g}_-^* .

We can now deduce that $\hat{\lambda}$ is Poisson, i.e. that for all $x \in T^*G/G$ and all $F, H \in \mathcal{F}(\mathfrak{g}_-^*)$

$$(407) \quad (\{F, H\}_{\mathfrak{g}^*} \circ \hat{\lambda})(x) = \{F \circ \hat{\lambda}, H \circ \hat{\lambda}\}_{T^*G/G}(x).$$

We just use (in order) the facts that:

- (i) π is surjective, so that for all $x \in T^*G/G$ there is an $\alpha_g \in T^*G$ with $x = \pi(\alpha_g) \equiv \text{Orb}(\alpha_g)$;
- (ii) $\mathbf{J}_R = \hat{\lambda} \circ \pi$;
- (iii) \mathbf{J}_R is Poisson; and
- (iv) π is Poisson:

$$(408) \quad (\{F, H\}_{\mathfrak{g}^*} \circ \hat{\lambda})(x) = \{F, H\}_{\mathfrak{g}^*} \circ (\hat{\lambda} \circ \pi)(\alpha_g)$$

$$(409) \quad = \{F, H\}_{\mathfrak{g}^*} \circ \mathbf{J}_R(\alpha_g) = \{F \circ \mathbf{J}_R, H \circ \mathbf{J}_R\}_{T^*G}(\alpha_g)$$

$$(410) = \{F \circ \hat{\lambda}, H \circ \hat{\lambda}\}_{T^*G/G}(\pi(\alpha_g)) \equiv \{F \circ \hat{\lambda}, H \circ \hat{\lambda}\}_{T^*G/G}(x) . .$$

■

7.3 Meshing with the symplectic structure on T^*G : invariant functions

We turn to giving more information about the situation described by the Lie-Poisson reduction theorem. The general idea will be that the Lie-Poisson bracket on \mathfrak{g}^* meshes with the canonical symplectic structure on T^*G . This will be made precise in two ways: the first is discussed in the first two Subsections, the second is discussed in the third Subsection.

The first discussion will have three stages:

- (i) we show that scalars on \mathfrak{g}^* , $F \in \mathcal{F}(\mathfrak{g}^*)$, are in one-one correspondence with scalars on T^*G that are constant on the orbits of the cotangent lift of left translation, which will be called *left-invariant* functions; and similarly, for the cotangent lift of right translation (a correspondence with *right-invariant* functions);
- (ii) we take the usual canonical Poisson bracket in T^*G of these left-invariant or right-invariant scalars; and restrict this bracket to \mathfrak{g}^* regarded as the cotangent space T_e^*G at the identity $e \in G$; and then
- (iii) we show that this restriction is the Lie-Poisson bracket on \mathfrak{g}^* : the familiar positive one for *right*-invariant functions, and the new negative one of eq. 402 for the *left*-invariant functions.

We do stages (i) and (ii) in Section 7.3.1. These stages will not involve the choice between the positive and negative Lie-Poisson brackets. But stage (iii), in Section 7.3.2, will involve this choice. It will be a one-liner corollary of Section 6.4.2's result that equivariant momentum maps are Poisson maps, eq. 365; (unsurprisingly, in that we also used this result in Section 7.2's proof of the reduction theorem).

In the third Subsection, we use invariant functions to show a different sense in which the Lie-Poisson bracket on \mathfrak{g}^* meshes with the symplectic structure on T^*G . Namely, we derive the Lie-Poisson bracket on \mathfrak{g}^* from the Poisson reduction theorem of Section 5.5, by using the ideas of invariant functions and momentum functions.

7.3.1 Left-invariant and right-invariant functions on T^*G

We say that a function $F : T^*G \rightarrow \mathbb{R}$ is *left-invariant* if for all $g \in G$, and all $\alpha_g \in T_g^*G$

$$(411) \quad (F \circ T^*L_g)(\alpha_g) = F(\alpha_g)$$

where T^*L_g is the cotangent lift of $L_g : G \rightarrow G$. Similarly, $F : T^*G \rightarrow \mathbb{R}$ is called *right-invariant* if for all $g \in G$

$$(412) \quad (F \circ T^*R_g) = F .$$

So if $F : T^*G \rightarrow \mathbb{R}$ is left-invariant or right-invariant, it is determined by its values for arguments in $T_e^*G = \mathfrak{g}^*$.

Since any $\alpha \in \mathfrak{g}^*$ is mapped by $T^*L_{g^{-1}} \equiv (T^*L_g)^{-1}$ to an element of T_g^*G , a function is left-invariant iff it is constant on the orbits of the various T^*L_g for $g \in G$, i.e. constant on the orbits of the cotangent lift of left translation. Similarly, a function is right-invariant iff it is constant on the orbits of the cotangent lift of right translation.

So left-invariant functions induce well-defined functions on the quotient space T^*G/G ; and so, by Section 7.2, on its diffeomorphic (indeed Poisson manifold) copy \mathfrak{g}^* . Similarly for right-invariant functions.

But let us for the moment consider the smooth left-invariant (or right-invariant) functions on T^*G , rather than the induced maps on the quotient space. We will denote the space of all smooth left-invariant functions on T^*G by $\mathcal{F}_L(T^*G)$, and similarly the space of smooth right-invariant functions by $\mathcal{F}_R(T^*G)$.

Recalling (from the discussion after eq. 120) that cotangent lifts are symplectic maps, i.e. T^*L_g and T^*R_g are symplectic maps on T^*G , it follows immediately that $\mathcal{F}_L(T^*G)$ and $\mathcal{F}_R(T^*G)$ are each closed under the canonical Poisson bracket on T^*G . So they are each a Lie algebra with this bracket.

Now we can use the momentum maps \mathbf{J}_L and \mathbf{J}_R of Example (3) of Section 6.5.3 to extend any scalar $F : \mathfrak{g}^* \rightarrow \mathbb{R}$, i.e. $F \in \mathcal{F}(\mathfrak{g}^*)$, to a left-invariant, or right-invariant, scalar on T^*G .

Thus, given $F : \mathfrak{g}^* \rightarrow \mathbb{R}$ and $\alpha_g \in T_g^*G$, we define $F_L \in \mathcal{F}_L(T^*G)$ by

$$(413) \quad F_L(\alpha_g) := (F \circ \mathbf{J}_R)(\alpha_g) \equiv (F \circ T_e^*L_g)(\alpha_g) .$$

So F_L is by construction left-invariant, and is called the *left-invariant extension* of F from \mathfrak{g}^* to T^*G .

One similarly defines the *right-invariant extension* $F_R \in \mathcal{F}_R(T^*G)$ of any $F \in \mathcal{F}(\mathfrak{g}^*)$ by

$$(414) \quad F_R(\alpha_g) := (F \circ \mathbf{J}_L)(\alpha_g) \equiv (F \circ T_e^* R_g)(\alpha_g) .$$

Then the maps

$$(415) \quad F \in \mathcal{F}(\mathfrak{g}^*) \mapsto F_L \in \mathcal{F}_L(T^*G) \quad \text{and} \quad F \in \mathcal{F}(\mathfrak{g}^*) \mapsto F_R \in \mathcal{F}_R(T^*G)$$

are vector space isomorphisms (exercise for the reader!) whose inverse is just restriction to the fiber $T_e^*G = \mathfrak{g}^*$.

This completes what we called ‘stages (i) and (ii)’: describing a correspondence between scalars on \mathfrak{g}^* and scalars on T^*G that are constant on the orbits of the cotangent lifts of left and right translation; and considering the canonical Poisson bracket (on T^*G) of these scalars, i.e. the Lie algebras $\mathcal{F}_L(T^*G)$ and $\mathcal{F}_R(T^*G)$.

7.3.2 Recovering the Lie-Poisson bracket

We now do stage (iii): we show that the restriction of the canonical Poisson bracket on T^*G of the right/left invariant functions, to \mathfrak{g}^* regarded as the cotangent space T_e^*G at the identity $e \in G$, is the positive/negative Lie-Poisson bracket.

Since the inverses of the maps eq. 415 are just restriction to the fiber $T_e^*G = \mathfrak{g}^*$, it suffices to show that the maps eq. 415 are Lie algebra isomorphisms. More precisely:

Recovery of the Lie-Poisson bracket Using the positive Lie-Poisson bracket on \mathfrak{g}^* (we write \mathfrak{g}_+^*): $F \mapsto F_R$ is a Lie algebra isomorphism. Similarly: using the negative Lie-Poisson bracket on \mathfrak{g}^* (we write \mathfrak{g}_-^*): $F \mapsto F_L$ is a Lie algebra isomorphism.

That is: for all $F, H \in \mathcal{F}(\mathfrak{g}^*)$

$$(416) \quad \{F, H\}_+ = \{F_R, H_R\}_{T^*G} \big|_{\mathfrak{g}^*} \quad ; \quad \{F, H\}_- = \{F_L, H_L\}_{T^*G} \big|_{\mathfrak{g}^*}$$

Proof. Consider $\mathbf{J}_L : T^*G \rightarrow \mathfrak{g}^* \equiv \mathfrak{g}_+^*$, $\mathbf{J}_L = T_e^* R_g$. \mathbf{J}_L is an equivariant momentum map. So, by the result eq. 365 of Section 6.4.2, it is Poisson. That is:

$$(417) \quad \{F, H\}_+ \circ \mathbf{J}_L = \{F \circ \mathbf{J}_L, H \circ \mathbf{J}_L\}_{T^*G} = \{F_R, H_R\}_{T^*G} .$$

Restricting eq. 417 to \mathfrak{g}^* gives the first equation of eq. 416.

Similarly, one proves the second equation by using the fact that $\mathbf{J}_R : T^*G \rightarrow \mathfrak{g}^* \equiv \mathfrak{g}_-^*$, $\mathbf{J}_R = T_e^* L_g$ is an equivariant momentum map and so is Poisson. That is:

$$(418) \quad \{F, H\}_- \circ \mathbf{J}_R = \{F \circ \mathbf{J}_R, H \circ \mathbf{J}_R\}_{T^*G} = \{F_L, H_L\}_{T^*G} .$$

We then restrict eq. 418 to \mathfrak{g}^* . ■

7.3.3 Deriving the Lie-Poisson bracket

Our discussion so far, in both Section 7.2 and the two previous Subsections, has taken the Lie-Poisson bracket (whether positive or negative) as given. We now show, using invariant functions and Section 6.5.1's idea of momentum functions, how to *derive* the Lie-Poisson bracket on \mathfrak{g}^* .

So this derivation will amount to another, more “constructive”, proof of the Lie-Poisson reduction theorem. As in Section 7.2's proof, two main ingredients will be:

- (a) the diffeomorphism $\hat{\lambda}$ between T^*G/G and \mathfrak{g}^* (eq. 229 or 404 or 405), and
- (b) the Poisson reduction theorem of Section 5.5, applied to G 's action on T^*G .

But instead of Section 7.2's proof's using the facts that (i) the momentum maps $\mathbf{J}_R \equiv T_e^*L_g$ and $\mathbf{J}_L \equiv T_e^*R_g$ are equivariant and (ii) equivariant momentum maps are Poisson, we will now use the ideas of invariant functions and momentum functions.

We begin by recalling that (since left translation is a diffeomorphism of G , and the cotangent lift of any diffeomorphism of a manifold to its cotangent bundle is symplectic), the Poisson reduction theorem implies that there is a unique Poisson structure on T^*G/G such that $\pi : T^*G \rightarrow T^*G/G$ is Poisson. We now use the diffeomorphism $\hat{\lambda} : T^*G/G \rightarrow \mathfrak{g}^*$ to transfer this Poisson structure to \mathfrak{g}^* . Let us call the result $\{, \}_-$. Though this is not to be read (yet!) as the negative Lie-Poisson bracket, our aim now is to calculate that it is in fact this bracket.

Notice first that since the momentum map $\mathbf{J}_R : T^*G \rightarrow \mathfrak{g}^*$, $\alpha_g \mapsto (T_e^*L_g)\alpha_g$ is equal to $\hat{\lambda} \circ \pi$ (eq. 406), we know that \mathbf{J}_R is Poisson with respect to this induced bracket on \mathfrak{g}^* . That is

$$(419) \quad \{F, H\}_- \circ \mathbf{J}_R(\alpha_g) = \{F \circ \mathbf{J}_R, H \circ \mathbf{J}_R\}_{T^*G}(\alpha_g) = \{F_L, H_L\}_{T^*G}(\alpha_g).$$

To calculate the right hand side, we will apply the ideas of invariant functions and momentum functions to each argument of the bracket; in particular to the first:

$$(420) \quad F_L(\alpha_g) = F(T_e^*L_g \cdot \alpha_g).$$

We observe that since a Poisson bracket depends only on the values of first derivatives, we can replace $F \in \mathcal{F}(\mathfrak{g}^*)$ by its linearization. That is, we can assume F is linear, so that at any point $\alpha \in \mathfrak{g}^*$, $F(\alpha) = \langle \alpha; \nabla F \rangle$, where ∇F is a constant in $\mathfrak{g} \equiv \mathfrak{g}^{**}$. Applying this, and the definition of a momentum function eq. 369, to eq. 420, we get:

$$(421) \quad F(T_e^*L_g \cdot \alpha_g) = \langle T_e^*L_g \cdot \alpha_g; \nabla F \rangle = \langle \alpha_g; T_eL_g \cdot \nabla F \rangle = \mathcal{P}(X_{\nabla F})(\alpha_g),$$

where the last equation applies the definition of a momentum function to the left-invariant vector field on G , $X_\xi(g) \equiv T_eL_g(\xi)$, for the case $\xi = \nabla F$.

Now we apply to eq. 421, in order: eq. 374, the definition of the Lie algebra bracket (cf. eq. 74), eq. 369 again, and the definition of left-invariant vector fields. We get:

$$(422) \quad \{F_L, H_L\}_{T^*G}(\alpha_g) = \{\mathcal{P}(X_{\nabla F}), \mathcal{P}(X_{\nabla H})\}_{T^*G}(\alpha_g) = -\mathcal{P}([X_{\nabla F}, X_{\nabla H}])(\alpha_g)$$

$$(423) \quad = -\mathcal{P}(X_{[\nabla F, \nabla H]})(\alpha_g) = -\langle \alpha_g; X_{[\nabla F, \nabla H]} \rangle$$

$$(424) = -\langle \alpha_g; T_e L_g([\nabla F, \nabla H]) \rangle = -\langle T_e^* L_g(\alpha_g); [\nabla F, \nabla H] \rangle .$$

Combining eq. 419 and eq. 424, and writing $\alpha \in \mathfrak{g}^*$ for $(T_e^* L_g)\alpha_g \equiv \mathbf{J}_R(\alpha_g)$, we have our result:

$$(425) \quad \{F, H\}_-(\alpha) = -\langle \alpha; [\nabla F, \nabla H] \rangle .$$

One similarly derives the positive Lie-Poisson bracket by considering right-invariant extensions of linear functions. The minus sign coming from eq. 374 is cancelled by the sign reversal in the Lie bracket of right-invariant vector fields. That is, it is cancelled by a minus sign coming from eq. 86.

7.4 Reduction of dynamics

We end our account of the Lie-Poisson reduction theorem by discussing the reduction of dynamics from T^*G to \mathfrak{g}^* .

We can be brief since we have already stated the main idea, when discussing the Poisson reduction theorem; cf. (2)(A) in Section 5.5. Thus recall that (under the conditions of the theorem) a G -invariant Hamiltonian function on a Poisson manifold M , $H : M \rightarrow \mathbb{R}$, defines a corresponding function h on M/G by $H = h \circ \pi$, where π is the projection $\pi : M \rightarrow M/G$; and since π is Poisson, and so pushes Hamiltonian flows forward to Hamiltonian flows, π pushes X_H on M to X_h on M/G :

$$(426) \quad T\pi \circ X_H = X_h \circ \pi .$$

Applying this, in particular eq. 426, to the Lie-Poisson reduction theorem, we get

Reduction of dynamics Let $H : T^*G \rightarrow \mathbb{R}$ be left-invariant. That is: the function $H^- := H|_{\mathfrak{g}^*}$ on \mathfrak{g}^* satisfies

$$(427) \quad H(\alpha_g) = H^-(\mathbf{J}_R(\alpha_g)) \equiv H^-(T_e^* L_g \cdot \alpha_g) , \quad \alpha_g \in T_g^* G .$$

Then \mathbf{J}_R pushes X_H forward to X_{H^-} . Or in terms of the flows $\phi(t)$ and $\phi^-(t)$ of X_H and X_{H^-} respectively:

$$(428) \quad \mathbf{J}_R(\phi(t)(\alpha_g)) = \phi^-(t)(\mathbf{J}_R(\alpha_g)) .$$

Similar statements hold for a right-invariant function $H : T^*G \rightarrow \mathbb{R}$, its restriction $H^+ := H|_{\mathfrak{g}^*}$ and $\mathbf{J}_L \equiv T_e^* R_g$.

Besides, we already know the vector field of H^- on \mathfrak{g}^* . For eq. 261 in (3) of Section 5.2.4 gave a basis-independent expression of Hamilton's equations on \mathfrak{g}^* in terms of ad^* . We just need to note that since we are now using the negative Lie-Poisson bracket on \mathfrak{g}^* , all terms in the deduction (eq. 261) apart from the left hand side, get a minus sign. So writing $\alpha \in \mathfrak{g}^*$, eq. 261 for the vector field X_{H^-} becomes:

$$(429) \quad \frac{d\alpha}{dt} = -ad_{\nabla H^-(\alpha)}^*(\alpha).$$

On the other hand, we can go in the other direction, reconstructing the dynamics on T^*G from eq. 429 on \mathfrak{g}^* . The statement of the main result, below, is intuitive, in that the “reconstruction equation” for $g(t) \in T^*G$ is

$$(430) \quad g^{-1}\dot{g} = \nabla H^- .$$

This is intuitive since it returns us to the basic idea of mechanics on \mathfrak{g} and \mathfrak{g}^* , viz. that the map

$$(431) \quad \lambda_g : \dot{g} \in T_g G \mapsto \lambda_g(\dot{g}) := (T_g L_{g^{-1}})\dot{g} \in \mathfrak{g}$$

maps the generalized velocity to its body representation; cf. eq. 205. However, the proof of this result is involved [Marsden and Ratiu, 1999, theorems 13.4.3, 13.4.4, p. 423–426]; so we only state the result. It is: —

Reconstruction of dynamics Suppose given a Lie group G , a left-invariant $H : T^*G \rightarrow \mathbb{R}$, its restriction $H^- := H|_{\mathfrak{g}^*}$, and an integral curve $\alpha(t)$ of the Lie-Poisson Hamilton's equations eq. 429 on \mathfrak{g}^* , with the initial condition $\alpha(0) = T_e^* L_{g_0}(\alpha_{g_0})$. Then the integral curve in T^*G of X_H is given by

$$(432) \quad T_{g(t)}^* L_{g(t)^{-1}}(\alpha(t)) ;$$

where $g(t)$ is the solution of the reconstruction equation

$$(433) \quad g^{-1}\dot{g} = \nabla H^-$$

with initial condition $g(0) = g_0$.

7.5 *Envoi: the Marsden-Weinstein-Meyer theorem*

I emphasize that our discussion of reduction has only scratched the surface: after all this Section has been relatively short! But now that the reader is armed with the long and leisurely exposition from Section 3 onwards, they are well placed to pursue the topic of reduction; e.g. through this Chapter's main sources, [Abraham and Marsden, 1978; Arnold, 1989; Olver, 2000; Marsden and Ratiu, 1999].

In particular, the reader can now relate the Lie-Poisson reduction theorem to another main theorem about symplectic reduction, usually called the Marsden–Weinstein–Meyer or Marsden–Weinstein theorem (after these authors' papers in

1973 and 1974).

This theorem concerns a symplectic action of a Lie group G on a symplectic manifold (M, ω) . For the sake of completeness, and to orient the reader to Landsman's discussion of this theorem (this vol., ch. 5, especially Section 4.5), it is worth stating it (as usual, for the finite-dimensional case only), together with the lemma used to prove it, and the ensuing reduction of dynamics. These statements will also round off our discussion by illustrating how some notions expounded from Section 3 onwards, but not used in this Section, are nevertheless useful — e.g. in stating the hypotheses of this theorem.

So suppose the Lie group G acts symplectically (eq. 311) on the symplectic manifold (M, ω) ; and that $\mathbf{J} : M \rightarrow \mathfrak{g}^*$ is an Ad^* -equivariant momentum map for this action (eq. 357 and 362). Assume also that $\alpha \in \mathfrak{g}^*$ is a regular value of \mathbf{J} , i.e. that at every point $x \in \mathbf{J}^{-1}(\alpha)$, $T_x \mathbf{J}$ is surjective. So the submersion theorem of (1) of Section 3.3.1 applies; in particular, $\mathbf{J}^{-1}(\alpha)$ is a sub-manifold of M with dimension $\dim(M) - \dim(\mathfrak{g}^*) \equiv \dim(M) - \dim(G)$.

Let G_α be the isotropy group (eq. 148) of α under the co-adjoint action, i.e.

$$(434) \quad G_\alpha := \{g \in G \mid Ad_{g^{-1}}^* \alpha = \alpha\}.$$

So since \mathbf{J} is Ad^* -equivariant under G_α , the quotient space $M_\alpha := \mathbf{J}^{-1}(\alpha)/G_\alpha$ is well-defined.

Now assume that G_α acts freely and properly on $\mathbf{J}^{-1}(\alpha)$, so that (Section 4.3.0.2) the quotient space $M_\alpha = \mathbf{J}^{-1}(\alpha)/G_\alpha$ is a manifold. M_α is the *reduced phase space* (corresponding to the momentum value α).

Now we assert:

Marsden-Weinstein-Meyer theorem M_α has a natural symplectic form ω_α induced from (M, ω) as follows. Let u, v be two vectors tangent to M_α at some point $p \in M_\alpha$: so p is an orbit of G_α 's action on $\mathbf{J}^{-1}(\alpha)$, and $u, v \in T_p M_\alpha$. Then u and v are obtained, respectively, from some vectors u' and v' tangent to $\mathbf{J}^{-1}(\alpha)$ at some point $x \in \mathbf{J}^{-1}(\alpha)$ of the orbit p , by the projection $\pi_\alpha : \mathbf{J}^{-1}(\alpha) \rightarrow M_\alpha$. That is:

$$(435) \quad T\pi_\alpha(u') = u \ ; \ T\pi_\alpha(v') = v \ .$$

It turns out that the value assigned by M 's symplectic form ω is the same whatever choice of x, u', v' is made. So we define the symplectic form ω_α on M_α as assigning this value. In other words: writing π_α for the projection, $i_\alpha : \mathbf{J}^{-1}(\alpha) \rightarrow M$ for the inclusion, and $*$ for pullback:

$$(436) \quad \pi_\alpha^* \omega_\alpha = i_\alpha^* \omega \ .$$

The proof of this theorem uses the following Lemma. Let us write $G \cdot x$ for the orbit $\text{Orb}(x)$ of x under the action of all of G , and similarly $G_\alpha \cdot x$ for the orbit under G_α , i.e. $\{\Phi(g, x) \mid g \in G_\alpha\}$. Then the Lemma states:

For any $x \in \mathbf{J}^{-1}(\alpha)$:

- (i): $T_x(G_\alpha \cdot x) = T_x(G \cdot x) \cap T_x(\mathbf{J}^{-1}(\alpha))$; and
- (ii): $T_x(G \cdot x)$ and $T_x(\mathbf{J}^{-1}(\alpha))$ are ω -orthogonal complements of one another in TM . That is: for all $u' \in T_xM$:
 $u' \in T_x(\mathbf{J}^{-1}(\alpha))$ iff $\omega(u', v') = 0$ for all $v' \in T_x(G \cdot x)$.

Both the Lemma and the theorem are each proven in some dozen lines. For details, cf. [Abraham and Marsden, 1978, Theorems 4.3.1–2, pp. 299–300; Arnold, 1989, Appendix 5.B, pp. 374–376].

Two final remarks. (1): The reduction of dynamics secured by the Marsden-Weinstein-Meyer theorem is similar to what we have seen before, for both the Poisson reduction theorem ((2) of Section 5.5), and the Lie-Poisson reduction theorem (Section 7.4). One proves, again in a few lines [Abraham and Marsden, 1978, Theorems 4.3.5, p. 304]:

Marsden-Weinstein-Meyer reduction of dynamics Let $H : M \rightarrow \mathbb{R}$ be invariant under the action of G on M , so that by Noether’s theorem for momentum maps (Section 6.2) \mathbf{J} is conserved, i.e. $\mathbf{J}^{-1}(\alpha)$ is invariant under the flow $\phi(t)$ of X_H on M . Then $\phi(t)$ commutes with the action of G_α on $\mathbf{J}^{-1}(\alpha)$ (i.e. $\phi(t) \circ \Phi_g = \Phi_g \circ \phi(t)$ for $g \in G_\alpha$), and so defines a flow $\hat{\phi}(t)$ on M_α such that $\pi_\alpha \circ \phi(t) = \hat{\phi}(t) \circ \pi_\alpha$, i.e.

$$(437) \quad \begin{array}{ccc} \mathbf{J}^{-1}(\alpha) & \xrightarrow{\pi_\alpha} & M_\alpha \\ \uparrow \phi(t) & & \uparrow \hat{\phi}(t) \\ \mathbf{J}^{-1}(\alpha) & \xrightarrow{\pi_\alpha} & M_\alpha \end{array}$$

The flow $\hat{\phi}(t)$ is Hamiltonian with the Hamiltonian H_α defined by $H_\alpha \circ \pi_\alpha = H \circ i_\alpha$.

(2): I said at the start of this Subsection that the reader can now relate the Lie-Poisson reduction theorem to the Marsden-Weinstein-Meyer theorem. It is not hard to show that the former is an example of the latter. As the symplectic manifold M one takes T^*G , acted on symplectically by the cotangent lift of left translation. So we know (from (3) of Section 6.5.3) that $\mathbf{J}_L := T_e^*R_g$ is an Ad^* -equivariant momentum map ... and so on: I leave this as an exercise for the reader! The answer is supplied at Arnold [1989, 377, 321] and Abraham and Marsden [1978, 302]. (Abraham and Marsden call it the ‘Kirillov-Kostant-Souriau theorem’.)

Suffice it to say here that this exercise gives another illustration of one of our central themes, that \mathfrak{g}^* ’s symplectic leaves are the orbits of the co-adjoint representation. For the reduced phase space M_α is naturally identifiable with the co-adjoint orbit $\text{Orb}(\alpha)$ of $\alpha \in \mathfrak{g}^*$, with the symplectic forms also naturally identified; (cf. also result (2) at the end of Section 5.4).

ACKNOWLEDGEMENTS

I am grateful to audiences in Irvine, Oxford, Princeton and Santa Barbara; to several colleagues for encouragement; and to Gordon Belot, Klaas Landsman, Barrie Tonkinson, David Wallace, and especially Graeme Segal, for helpful, and patient!, conversations, correspondence and corrections.

BIBLIOGRAPHY

- [Abraham and Marsden, 1978] R. Abraham and J. Marsden. *Foundations of Mechanics*, second edition: Addison-Wesley, 1978.
- [Arnold, 1973] V. Arnold. *Ordinary Differential Equations*, MIT Press, 1973.
- [Arnold, 1989] V. Arnold. *Mathematical Methods of Classical Mechanics*, Springer, 1989 (second edition).
- [Barbour and Bertotti, 1982] J. Barbour and B. Bertotti. Mach's principle and the structure of dynamical theories, *Proceedings of the Royal Society of London A* **382**, p. 295-306, 1982.
- [Belot, 1999] G. Belot. Relationism rehabilitated, *International Studies in Philosophy of Science* **13**, p. 35-52, 1999.
- [Belot, 2000] G. Belot. Geometry and motion, *British Journal for the Philosophy of Science* **51**, p. 561-596, 2000.
- [Belot, 2001] G. Belot. The principle of sufficient reason', *Journal of Philosophy* **98**, p. 55-74, 2001.
- [Belot, 2003] G. Belot. Notes on symmetries, in [Brading and Castellani, 2003, pp. 393-412].
- [Belot, 2003a] G. Belot. Symmetry and gauge freedom, *Studies in the History and Philosophy of Modern Physics* **34**, p. 189-225, 2003.
- [Belot, 2006] G. Belot (this volume), 2006.
- [Belot and Earman, 2001] G. Belot and J. Earman. Pre-Socratic quantum gravity, in C. Callender and N. Huggett, eds. *Physics meets Philosophy at the Planck Scale*, Cambridge University Press, pp. 213-255, 2001.
- [Bishop and Goldberg, 1980] R. Bishop and S. Goldberg. *Tensor Analysis on Manifolds*, New York: Dover, 1980.
- [Boudri, 2002] J. Boudri. *What was Mechanical about Mechanics: the Concept of Force between Metaphysics and Mechanics from Newton to Lagrange*, Dordrecht: Kluwer Academic, 2002.
- [Brading and Castellani, 2003] K. Brading and E. Castellani, eds. *Symmetry in Physics*, Cambridge University Press, 2003.
- [Brading and Castellani, 2006] K. Brading and E. Castellani (this volume), 1006.
- [Butterfield, 2004] J. Butterfield. Some Aspects of Modality in Analytical mechanics', in *Formal Teleology and Causality*, ed. M. Stöltzner, P. Weingartner, Paderborn: Mentis, 2004. Available at Los Alamos arXiv: <http://arxiv.org/abs/physics/0210081> ; and at Pittsburgh archive: <http://philsci-archive.pitt.edu/archive/00001192>.
- [Butterfield, 2004a] J. Butterfield. Between Laws and Models: Some Philosophical Morals of Lagrangian Mechanics, 2004; available at Los Alamos arXiv: <http://arxiv.org/abs/physics/0409030> ; and at Pittsburgh archive: <http://philsci-archive.pitt.edu/archive/00001937/>.
- [Butterfield, 2005] J. Butterfield. Between Laws and Models: Some Philosophical Morals of Hamiltonian Mechanics, in preparation, 2005.
- [Butterfield, 2006] J. Butterfield. On Symmetry and Conserved Quantities in Classical Mechanics, forthcoming in a *Festschrift* for Jeffrey Bub, ed. W. Demopoulos and I. Pitowsky, Kluwer: University of Western Ontario Series in Philosophy of Science, 2006. available at Los Alamos arXiv: <http://arxiv.org/abs/physics/> ; and at Pittsburgh archive: <http://philsci-archive.pitt.edu/archive/00002362/>
- [Darling, 1994] R. Darling. *Differential Forms and Connections*, Cambridge University Press, 1994.
- [Dickson, 2006] M. Dickson (this volume), 2006.
- [Desloge, 1982] E. Desloge. *Classical Mechanics*, John Wiley, 1982.

- [Earman, 2003] J. Earman. Tracking down gauge: an ode to the constrained Hamiltonian formalism, in [Brading and Castellani, 2003, pp. 140-162].
- [Goldstein, 1950] H. Goldstein. *Classical Mechanics*, Addison-Wesley, 1950; (1966 third printing).
- [Grattan-Guinness, 2000] I. Grattan-Guinness. A sideways look at Hilbert's twenty-three problems of 1900, *Notices of the American Mathematical Society* **47**, p. 752-757, 2000.
- [Grattan-Guinness, 2006] I. Grattan-Guinness. Classical mechanics as a formal(ised) science, in B. Loewe (ed.), *Foundations of the Formal Sciences*, Kluwer, 2006.
- [Hawkins, 2000] T. Hawkins. *Emergence of the Theory of Lie Groups: an essay in the history of mathematics 1869-1926*, New York: Springer, 2000.
- [Jammer, 1957] M. Jammer. *Concepts of Force*, Harvard University Press, 1957.
- [Jammer, 1961] M. Jammer. *Concepts of Mass in Classical and Modern Physics*, Harvard University Press, 1961; republished by Dover in 1997.
- [Jammer, 2000] M. Jammer. *Concepts of Mass in Contemporary Physics and Philosophy*, Princeton University Press, 2000.
- [Johns, 2005] O. Johns. *Analytical Mechanics for Relativity and Quantum Mechanics*, Oxford University Press, 2005.
- [Kirillov, 1976] A. Kirillov. *Elements of the Theory of Representations*, Grunlehen Math. Wiss., Springer-Verlag, 1976.
- [Kragh, 1999] H. Kragh. *Quantum Generations*, Princeton University Press, 1999.
- [Landsman, 2006] N. Landsman (this volume), Between Classical and Quantum.
- [Lie, 1890] S. Lie. *Theorie der Transformationsgruppen: zweiter abschnitt*, Leipzig: B.G.Teubner, 1890.
- [Lutzen, 1995] J. Lutzen. *Denouncing Forces; Geometrizing Mechanics: Hertz's Principles of Mechanics*, Copenhagen University Mathematical Institute Preprint Series No 22, 1995
- [Lutzen, 2003] J. Lutzen. Between rigor and applications: developments in the concept of function in mathematical analysis, in *Cambridge History of Science, vol. 5: The modern physical and mathematical sciences*, ed. M.J. Nye, p. 468-487, 2003.
- [Lutzen, 2005] J. Lutzen. *Mechanistic Images in Geometric Form: Heinrich Hertz's 'Principles of Mechanics'*, Oxford University Press, 2005.
- [Marsden and Hughes, 1982] J. Marsden and T. Hughes. *Mathematical Foundations of Elasticity*, Prentice-Hall; Dover, (1982) 1994.
- [Marsden and Ratiu, 1999] J. Marsden and T. Ratiu. *Introduction to Mechanics and Symmetry*, second edition: Springer-Verlag, 1999.
- [McMullin, 1978] E. McMullin. *Newton on Matter and Activity*, University of Notre Dame Press, 1978.
- [Newton, 1687] I. Newton. *Principia: Mathematical Principles of Natural Philosophy*, 1687. Trans. I.B. Cohen and A. Whitman, Cambridge University Press, 1999.
- [Olver, 2000] P. Olver. *Applications of Lie Groups to Differential Equations*, second edition: Springer-Verlag, 2000.
- [Rovelli, 2006] C. Rovelli (this volume), Quantum Gravity.
- [Ryckman, 2005] T. Ryckman. *The Reign of Relativity: Philosophy in Physics 1915-1925* Oxford University Press, 2005.
- [Singer, 2001] S. Singer. *Symmetry in Mechanics: a Gentle Modern Introduction*, Boston: Birkhauser, 2001.
- [Slovik, 2002] E. Slovik. *Cartesian Spacetime: Descartes' Physics and the Relational Theory of Space and Motion*, Dordrecht: Kluwer Academic, 2002.
- [Stöltzner, 2003] M. Stöltzner. The Principle of Least Action as the Logical Empiricist's Shibboleth, *Studies in History and Philosophy of Modern Physics* **34B**, p. 285-318, 2003.
- [Wallace, 2003] D. Wallace. Time-dependent Symmetries: the link between gauge symmetries and indeterminism, in [Brading and Castellani, 2003, pp. 163-173].

This page intentionally left blank

THE REPRESENTATION OF TIME AND CHANGE IN MECHANICS

Gordon Belot

If time is objective the physicist must have discovered that fact, if there is Becoming the physicist must know it; but if time is merely subjective and Being is timeless, the physicist must have been able to ignore time in his construction of reality and describe the world without the help of time. . . . If there is a solution to the philosophical problem of time, it is written down in the equations of mathematical physics.

Perhaps it would be better to say that the solution is to be read between the lines of the physicist's writings. Physical equations formulate specific laws . . . but philosophical analysis is concerned with statements about the equations rather than with the content of the equations themselves.

— Reichenbach.¹

For many years I have been tormented by the certainty that the most extraordinary discoveries await us in the sphere of Time. We know less about time than about anything else.

— Tarkovsky.²

1 INTRODUCTION

This chapter is concerned with the representation of time and change in classical (i.e., non-quantum) physical theories. One of the main goals of the chapter is to attempt to clarify the nature and scope of the so-called problem of time: a knot of technical and interpretative problems that appear to stand in the way of attempts to quantize general relativity, and which have their roots in the general covariance of that theory.

The most natural approach to these questions is via consideration of more clear cases. So much of the chapter is given over to a discussion of the representation of time and change in other, better understood theories, starting with the most straightforward cases and proceeding through a consideration of cases that prepare one, in one sense or another, for the features of general relativity that are responsible for the problem of time.

¹[Reichenbach, 1991, 16 f.].

²[Tarkovsky, 1991, 53].

Let me begin by saying a bit about what sort of thing I have in mind in speaking of the representation of time and change in physical theories, grounding the discussion in the most tractable case of all, Newtonian physics.

As a perfectly general matter, many questions and claims about the content of a physical theory admit of two construals — as questions about structural features of solutions to the equations of motion of the theory, or as questions about structural features of these equations. For instance, on the one hand time appears as an aspect of the spacetimes in which physics unfolds — that is, as an aspect of the background in which the solutions to the equations of the theory are set. On the other, time is represented via its role in the laws of physics — in particular, in its role in the differential equations encoding these laws. So questions and claims about the nature of time in physical theories will admit of two sorts of reading.

Consider, for instance, the claim that time is homogeneous in Newtonian physics (or, as Newton would put it, that time flows equably). There are two sorts of fact that we might look to as grounding this claim.

1. There is a sense in which time is a separable aspect of the spacetime of Newtonian physics and there is a sense in which time, so considered, is homogenous.³
2. The laws of the fundamental-looking theories of classical mechanics (e.g., Newton's theory of gravity) are time translation invariant — the differential equations of these theories do not change their form when the origin of the temporal coordinate is changed — so the laws of such theories are indifferent to the identity of the instants of time.

In the Newtonian setting, these two sorts of considerations mesh nicely and provide mutual support: there is a consilience between the symmetries of the laws and the symmetries of spacetime. But in principle, the two sorts of consideration need not lead to the same sort of answer: one might consider a system in Newtonian spacetime that is subject to time-dependent forces; or one could set the Newtonian n -body problem in a spacetime which featured a preferred instant, but otherwise had the structure of Newtonian spacetime. And as one moves away from the familiar setting of Newtonian physics, it becomes even more important to distinguish the two approaches: in general relativity, the laws have an enormous (indeed, infinite-dimensional) group of symmetries while generic solutions have no symmetries whatsoever.

In discussing the representation of time and change, this chapter will focus on structural features of the laws of physical theories rather than on features of particular solutions. To emphasize this point, I will say that I am interested in the structure of this or that theory as a *dynamical theory*.

³(Neo)Newtonian spacetime is partitioned in a natural way by instants of absolute simultaneity, and time can be identified with the structure that the set of these instants inherits from the structure of spacetime: time then has the structure of an affine space modelled on the real numbers — so for any two instants, there is a temporal symmetry which maps one to the other.

I will approach my topics via the Lagrangian and Hamiltonian approaches to classical theories, two great over-arching — and intimately related — frameworks in which such topics are naturally addressed.⁴ Roughly speaking, in each of these approaches the content of the equations of a theory is encoded in certain structures on a space of possibilities associated with the theory.⁵ In the Lagrangian approach the featured space is the space of solutions to the equations of the theory, which for heuristic purposes we can identify with the space of possible worlds allowed by the theory.⁶ On the Hamiltonian side, the featured space is the space of initial data for the equations of the theory, which we can in the same spirit identify with the space of possible instantaneous states allowed by the theory.

In Newtonian mechanics, the reflection within the Lagrangian framework of the time translation invariance of the laws is that the space of solutions is itself invariant under time translations: given a set of particle trajectories in spacetime obeying Newton's laws of motion, we can construct the set of particle trajectories that result if all events are translated in time by amount t ; the latter set is a solution (i.e., is permitted by the laws of motion) if and only if the former set is; furthermore, the map that carries us from a solution to its time translate preserves the structure on the space of solutions that encodes the dynamics of the theory. Within the Hamiltonian framework, on the other hand, the time translation invariance of the laws is reflected by the existence of a map that sends an initial data set to the state it will evolve into in t units of time; again, this map leaves invariant the structure on the space that encodes the dynamics of the theory. So the temporal symmetry of the dynamics of the theory is reflected on the Lagrangian side by a notion of time translation and on the Hamiltonian side by a notion of time evolution.

The representation of change in Newtonian physics also takes different (but closely related) forms within the Lagrangian and Hamiltonian frameworks. Change consists in a system having different and incompatible properties at different times. We want to say, for instance, that there is a change in the observable properties of a two-body system if and only if the relative distance between the particles is different at different times.

⁴Why pursue our question within the realm of Lagrangian and Hamiltonian mechanics rather than working directly with the differential equations of theories? Because the benefits are large: these overarching approaches provide powerful mathematical frameworks in which to compare theories. And because the costs are minimal: almost every theory of interest can be put into Lagrangian or Hamiltonian form, without any obvious change of content. And because it leads us where we want to go: current attempts to understand the content of classical physical theories are necessarily shaped by efforts to construct or understand deeper, quantum theories; and it appears that a classical theory *must* be placed in Lagrangian or Hamiltonian form in order to be quantized.

⁵See remark 2 for the sense in which talk of possibilities is intended here.

⁶In the context of ordinary classical mechanics, one often thinks of Lagrangian mechanics as set in the velocity phase space — and thus as more closely associated with a space of initial data rather than with a space of solutions. However, this familiar approach presupposes an absolute notion of simultaneity, and for this reason it is usually dropped in favour of a spacetime covariant Lagrangian approach (in which the space of initial data plays no role) when one turns to relativistic theories. This is the point of view adopted below.

Hamiltonian Approach. Specifying the instantaneous dynamical state of such a system suffices to specify the instantaneous relative distance between the particles. So there is a function on the space of initial data corresponding to this quantity. A history of the system is a trajectory through the space of initial data. In our simple example, observable change occurs during a given history if and only if the function corresponding to the relative distance between the particles takes on different values at different points on the trajectory in question. More generally, in any Newtonian system, any quantity of physical interest (observable or not) is represented by a function on the space of initial data, and a trajectory in this space represents such quantities as changing if the corresponding functions take on different values at different points on the trajectory.

Lagrangian Approach. Clearly, no function on the space of solutions can represent a changeable quantity in the same direct way that functions on the space of initial data can. But for each t , there is a function on the space of solutions of our two-body problem that assigns to each solution the relative distance between the particles at time t according to that solution. Letting t vary, we construct a one-parameter family of functions on the space of solutions. A solution to the equations of motion represents the relative distance between the particles as changing if and only if different members of this one-parameter family of functions take on different values when evaluated on the given solution. And so on more generally: any changeable physical quantity corresponds to such a one-parameter family of functions on the space of solutions, and change is understood as in the simple two-body example.

So much for the sort of thing I have in mind in speaking of the representation of time and change in a physical theory. Before sketching the path that this chapter takes in discussing these topics, it will perhaps be helpful to say a bit about its ultimate goal — the clarification of the nature of the so-called problem of time. Discussions of the problem of time typically focus on Hamiltonian versions of general relativity, in which the focus is on the space of possible instantaneous geometries (metrics and second fundamental forms on Cauchy surfaces). This is somewhat unfortunate, since such approaches require from the start a division of spacetime into a family of spacelike hypersurfaces — which appears to be against the spirit of the usual understanding of the general covariance of the theory. In light of this fact, there is room for worry that some aspects of the problem of time as usually presented are consequences of this rather awkward way of proceeding. I take a somewhat different path, always anchoring my discussion in the Lagrangian approach, which takes as fundamental complete histories of systems rather than instantaneous states.

The view developed below is that, *roughly speaking*, the core of the problem of time is that in general relativity, when understood dynamically, there is no way to view time evolution or time translation as symmetries of the theory and, relatedly, there is no natural way to model change via functions on the spaces arising within

the Lagrangian and Hamiltonian approaches.⁷ This marks a respect in which general relativity, so conceived, is very different from preceding fundamental-looking theories.

The problem of time may sound — not very pressing. To be sure, there are puzzles here. Why should general relativity differ in this way from its predecessors? In predecessors to general relativity, the representation of time and the representation of change are tied together in a very neat package — what does the general relativistic replacement for this package look like? These are interesting questions. But then of course no one should expect time to be represented in general relativity as in its predecessors — that it presents an utterly new picture of time and space is one of the glories of the theory. And one might also think: since the structure of spacetime varies from solution to solution in general relativity, it is surely more appropriate to look at the representation of time in this or that physically realistic solution, rather than in the equations of the theory, if we want to understand what the theory is telling us about the nature of time in our world.

The problem of time assumes a more pressing aspect, however, when one considers the quantization of general relativity (or of any other theory that is generally covariant in the relevant sense). The project of constructing successor theories naturally focuses our attention on structural features of the theories at hand — in constructing successors, one is in the business of laying bets as to which such features of current theories will live on (perhaps in a new form), and which ones will be left behind. And known techniques of quantization require as input not just differential equations, but theories cast in Hamiltonian or Lagrangian form. So for those interested in quantizing general relativity, questions about the structure of the theory *qua* dynamical theory naturally loom large. And lacking solutions to the puzzles mentioned above, one expects conceptual difficulties in formulating (or extracting predictions from) any quantization of general relativity. So from this perspective, the problem of time is in fact quite pressing.

This chapter takes long route to the problem of time. I begin in section 2 with the briefest of introductions to Hamiltonian and Lagrangian mechanics, by way of motivating some of what follows. In section 3, I sketch some important concepts and results of symplectic geometry, the field of mathematics that underlies classical

⁷This formulation above only gives us a first approximation, for several reasons (each of which will be discussed more fully in following sections). (i) The problem of time only arises in those versions of general relativity most appropriate to the cosmological setting; in other applications of the theory, time is represented in a fashion very similar to that in which it is represented in special relativistic physics. (ii) In the treatment of ordinary time-dependent systems, time evolution and time translation are not symmetries of the theory — but this does not lead to any real problem in representing time and change in such theories, because one still has group actions that implement time evolution and time translation, even though they are not symmetries of the laws, and these suffice to erect an account of change very similar to that occurring in ordinary time-independent theories. (iii) In theories in which solutions are not defined globally in time, time evolution and time translation will not be implemented by group actions, but merely by local flows (these can be thought of as infinitesimal surrogates for group actions); these suffice for the construction of the familiar picture of change in such theories; but even these are absent in general relativity.

mechanics. The concepts introduced here are crucial for what follows: for well-behaved theories, the space of solutions (on the Lagrangian side) and the space of initial data (on the Hamiltonian side) both have symplectic structures. And we will see that various symplectic (or nearly symplectic) spaces arise even when one strays away from the ideal case. In section 4, I sketch the very powerful framework of modern Lagrangian mechanics, with its apparatus of local conservation laws.

In section 5, I sketch the Lagrangian and Hamiltonian pictures for ideally well-behaved theories satisfying the following conditions: (i) the background spacetime geometry admits a group of time translations and the Lagrangian of the theory is invariant (in a suitable sense) under the action of this group; (ii) specifying initial data for the equations of the theory suffices to determine a single maximal solution; (iii) this maximal solution is defined for all values of the time parameter. When these conditions hold, we find that there is a group of time translation symmetries operating on the space of solutions on the Lagrangian side, while on the Hamiltonian side there is a group implementing time evolution on the space of initial data. These two spaces are isomorphic, and the two group actions intertwine in a satisfying fashion. One is able to give a straightforward and appealing account of the way in which change is represented on either of these two fundamental spaces.

In section 6, I turn to the complications that must be introduced into the picture when one drops any one of the conditions (i)–(iii) of the preceding paragraph. Finally in section 7, I address the representation of time and change in general relativity. This leads directly to the problem of time.

As this outline makes clear, much of the chapter is given over to exposition of technical material. In order to keep the length reasonable, I have had to presume that the reader comes to this chapter with quite a bit of technical background. I have tried to write for an ideal reader who has previously studied general relativity or gauge theory, and hence feels comfortable with the basic concepts, results, and constructions of differential geometry (although at a few strategic points I have included discussion meant to jog the memory of such readers).

This chapter is founded upon the modern geometrical approach to Lagrangian mechanics that is presented in the barest sketch in section 4. This approach, developed relatively recently by mathematicians, provides a highly abstract framework for thinking about physical theories rather than a fully rigorous treatment of any given theory. It exists at the formal, differential-geometric level: the focus is on the geometric structure of various spaces and on the geometric content of equations and constructions; functional analytic details are held in abeyance. Much of the material sketched in other sections functions at this same level.

In content, this chapter overlaps somewhat with [Malament, this volume], [Rovelli, this volume], and [Brading and Castellani, this volume]. But it is most closely related to [Butterfield, this volume]. Butterfield's chapter provides a philosophical introduction to modern geometric approaches to mechanics; the present chapter is meant as an example of the application of this approach to a philosophical problem. The present chapter is, however, meant to be self-contained. And there is in

fact a considerable difference in emphasis between this chapter and Butterfield's: the latter is restricted to finite-dimensional systems, and focuses on the Hamiltonian side of things; the present chapter is primarily concerned with field theories, and focuses to a much greater extent on the Lagrangian approach.

REMARK 1 (Notation and Terminology). Elements of and structures on the space of solutions of a field theory are always indicated by capital letters (Greek or Latin) while elements of and structures on the space initial data of a field theory are always indicated by lower case letters (Greek or Latin). Boldface indicates three-vectors or three-vector-valued functions. In this chapter, a curve is officially a map from intervals of real numbers into a space that is a manifold or a mild generalization of a manifold — sometimes for emphasis I redundantly call a curve a *parameterized curve*. An *affinely parameterized curve* is a equivalence class of such curves, where two curves count as equivalent if they have the same image and their parameterization agrees up to a choice of origin.⁸ A *unparameterized curve* is an equivalence class of curves, under the equivalence relation where curves count as equivalent if they have the same image. I sometimes conflate a curve and its image.

REMARK 2 (Possible Worlds Talk). Below, especially in section 7, I sometimes speak of points of the space of solutions (initial data) as representing possible worlds (possible instantaneous states) permitted by the theory, even though I do not pretend to be involved in fine-grained matters of interpretation here. This sort of thing is meant only in a rough and heuristic way. The idea is that in trying to understand a theory, we are in part engaged in a search for a perspicuous formulation of the theory; and it is reasonable to hope that if a formulation is perspicuous, then there will exist a *prima facie* attractive interpretation of the theory according to which there is a bijection between the space of solutions (initial data) and the space of possible worlds (possible instantaneous states) admitted by the theory under that interpretation. This is not to deny that there may be reasons for ultimately rejecting such interpretations: a Leibnizean might settle on a standard formulation of classical mechanics, even though that means viewing the representation relation between solutions and possible worlds as many-to-one in virtue of the fact that solutions related by a time translation must be seen as corresponding to the same possible world.

2 HAMILTONIAN AND LAGRANGIAN MECHANICS

This section contains a *very* brief sketch of the Hamiltonian and Lagrangian approaches to the Newtonian n -body problem.⁹ The intended purpose is to motivate

⁸That is, an affinely parameterized curve is an equivalence class of curves under the equivalence relation according to which curves $\gamma_1 : [a, b] \rightarrow M$ and $\gamma_2 : [a, b] \rightarrow M$ are equivalent if and only if there exists $s \in \mathbb{R}$ such that $\gamma_1(t) = \gamma_2(t + s)$ for all $t \in [a, b]$.

⁹For textbooks approaching classical mechanics in a variety of styles, see, e.g., [Goldstein, 1953], [Lanczos, 1986], [Singer, 2001], [Marsden and Ratiu, 1994], [Arnold, 1989], [Arnold *et al.*,

some of what follows in later sections.

2.1 The n -Body Problem

We consider n gravitating point-particles. Let the mass of the i th particle be m_i . Working relative to a fixed inertial frame we write: $q := (\mathbf{q}_1, \dots, \mathbf{q}_n) = (q_1, \dots, q_{3n})$ for the positions of the particles, $\dot{q} := (\dot{\mathbf{q}}_1, \dots, \dot{\mathbf{q}}_n) = (\dot{q}_1, \dots, \dot{q}_{3n})$ for their velocities, and $\ddot{q} := (\ddot{\mathbf{q}}_1, \dots, \ddot{\mathbf{q}}_n) = (\ddot{q}_1, \dots, \ddot{q}_{3n})$ for their accelerations (in this chapter, boldface always indicates a three-vector). The gravitational force exerted on the i th particle by the j th particle is

$$\mathbf{F}_{ij} = \frac{m_i m_j}{r_{ij}^2} \mathbf{u}_{ij}, \quad (1)$$

where r_{ij} is the distance between the i th and j th particles, \mathbf{u}_{ij} is the unit vector pointing from the i th to the j th particle, and units have been chosen so that Newton's constant is unity. Of course, equation 1 is not well-defined for $r_{ij} = 0$. So from now on we assume that $q \in Q := \mathbb{R}^{3n}/\Delta$, where Δ is the *collision set* $\{q \in \mathbb{R}^{3n} : \mathbf{q}_i = \mathbf{q}_j \text{ for some } i \neq j\}$.

The net force acting on the i th particle is

$$\mathbf{F}_i = \sum_{j \neq i} \mathbf{F}_{ij}.$$

So the *equations of motion* for our theory are: $\mathbf{F}_i = m_i \ddot{\mathbf{q}}_i$.¹⁰ Resolving each force and acceleration vector into its components, we have $3n$ second-order differential equations. Roughly speaking, these equations have a well-posed initial value problem: specifying $3n$ values for the initial positions of our particles and $3n$ values for their initial velocities (momenta) determines a unique analytic solution to the equations of motion, which tells us what the positions and velocities (momenta) of the particles are at all other times at which these quantities are defined.¹¹

2.2 The Hamiltonian Approach

The basic variables of the Hamiltonian approach are the positions of the particles and the corresponding momenta, $p := (m_1 \dot{\mathbf{q}}_1, \dots, m_n \dot{\mathbf{q}}_n) = (m_1 q_1, \dots, m_n q_{3n})$. A state of the system, (q, p) , is specified by specifying the position and momentum of each particle. To each state we can assign a *kinetic energy*

$$T(q, p) := \sum_{i=1}^n \frac{1}{2m_i} |\mathbf{p}_i|^2$$

1997], and [Abraham and Marsden, 1978].

¹⁰This should be read as a differential equation constraining the allowed trajectories $q(t)$. Similarly for the other differential equations appearing in this section.

¹¹Note that some solutions fail to be defined for all values of t ; see example 33 below for discussion.

and a *potential energy*,

$$V(q, p) := \sum_{i < j} \frac{m_i m_j}{r_{ij}}.$$

Note that $\mathbf{F}_i = -\nabla_i V(q)$, where ∇_i is the gradient operator $(\frac{\partial}{\partial q_{3i-2}}, \frac{\partial}{\partial q_{3i-1}}, \frac{\partial}{\partial q_{3i}})$ corresponding to the configuration variables of the i th particle. So the potential energy encodes information about gravitational forces, while the kinetic energy can be thought of as encoding information about the inertial structure of Newtonian spacetime. So one might hope that together these quantities encode all of the physics of the n -body problem. This is indeed the case.

We introduce the space of initial data for the theory, $\mathcal{I} := \{(q, p): q \in Q\}$ and the Hamiltonian $H : \mathcal{I} \rightarrow \mathbb{R}$, $H(q, p) := T(q, p) + V(q)$. The Hamiltonian is thus just the total energy.

The original equation of motion $m_i \ddot{\mathbf{q}}_i = \mathbf{F}_i$ can be rewritten as $\dot{\mathbf{p}}_i = -\nabla_i V(q)$; or, since $\nabla_i T = 0$, as $\dot{\mathbf{p}}_i = -\nabla_i H$. In another notation, this becomes $\dot{p}_i = -\frac{\partial H}{\partial q_i}$. Furthermore, since the only term in H depending on p_i is of the form $\frac{1}{2m} p_i^2$, we find that $\frac{\partial H}{\partial p_i} = \dot{q}_i$.

In this way, we move from the original Newtonian equations to *Hamilton's equations*:

$$\dot{q}_i = \frac{\partial H}{\partial p_i} \quad \text{and} \quad \dot{p}_i = -\frac{\partial H}{\partial q_i} \quad (i = 1, \dots, 3n).$$

In fact, the ordinary Newtonian equations are equivalent to Hamilton's equations. So we see that the function $H = T + V$ encodes all of the dynamical content of the n -body problem.

Our present interest is in the geometry implicit in Hamilton's equations. Hamilton's equations gives us values for $\dot{q}_i(q, p)$ and $\dot{p}_i(q, p)$ at each point $(q, p) \in \mathcal{I}$. That is, Hamilton's equations give us a component expression for a tangent vector $X_H(q, p)$ at each point $(q, p) \in \mathcal{I}$. The vector field X_H on \mathcal{I} encodes the dynamics of our theory: through each point $(q_0, p_0) \in \mathcal{I}$ there is exactly one curve $(q(t), p(t)) : \mathbb{R} \rightarrow \mathcal{I}$ such that: (i) $(q(0), p(0)) = (q_0, p_0)$; and (ii) for each s , the tangent vector to the curve $(q(t), p(t))$ at $t = s$ is given by $X_H(q(s), p(s))$. This curve tells us that if the system is in state (q_0, p_0) at time $t = 0$, then it is in state $(q(s), p(s))$ at time $t = s$.

We can rewrite Hamilton's equations as:

$$(\dot{q}_1, \dots, \dot{q}_{3n}, \dot{p}_1, \dots, \dot{p}_{3n}) \begin{vmatrix} 0 & I \\ -I & 0 \end{vmatrix} = \left(\frac{\partial H}{\partial q_1}, \dots, \frac{\partial H}{\partial q_{3n}}, \frac{\partial H}{\partial p_1}, \dots, \frac{\partial H}{\partial p_{3n}} \right),$$

where I is the $3n \times 3n$ identity matrix. On the left-hand side we have a vector multiplied by a matrix; on the right-hand side another vector. Thinking of \mathcal{I} as a manifold, we can recognize the coordinate-independent objects standing behind this equation: on the left we have the tangent vector field X_H contracted with a

two-form; on the right, the differential dH (i.e., the exterior derivative of H). So we can re-write Hamilton's equations in a coordinate-independent form as:

$$\omega(X_H, \cdot) = dH,$$

where ω is the two-form on \mathcal{I} that assumes the form $\sum_i dq_i \wedge dp_i$ in our coordinates.

ω is a *symplectic form* on \mathcal{I} : a closed, nondegenerate two-form.¹² ω can be thought of as being somewhat like an anti-symmetric metric on \mathcal{I} (e.g., both sorts of object establish a preferred isomorphism between vector fields and one-forms). But that analogy cannot be taken too seriously in light of the following striking differences between the two sorts of objects:

1. The isometry group of a finite-dimensional Riemannian manifold is always finite-dimensional. But our symplectic form is invariant under an infinite-dimensional family of diffeomorphisms from \mathcal{I} to itself. We can see this as follows. Let us think of \mathcal{I} as the cotangent bundle of Q ; that is, we think of a point (q, p) as consisting of a point $q \in Q$ and a covector $p \in T_q^*Q$.¹³ A *cotangent coordinate system* on $\mathcal{I} = T^*Q$ arises as follows: choose arbitrary coordinates $\{q_i\}$ on Q and write $p \in T_q^*Q$ as $p = \sum p_i dq^i$, so that $\{q_i, p_j\}$ forms a set of coordinates on T^*Q . In any cotangent coordinate system,

$$\omega = \begin{vmatrix} 0 & I \\ -I & 0 \end{vmatrix}. \quad (2)$$

So ω is invariant under the transformation that carries us from one set of cotangent coordinates on \mathcal{I} to another. And the set of such transformations is infinite-dimensional, since any diffeomorphism $d : Q \rightarrow Q$ generates such a transformation.

2. One does not expect any manifold or bundle to carry a natural Riemannian metric. But if M is any finite-dimensional manifold, the cotangent bundle T^*M carries a canonical symplectic form, ω , that takes the form $\omega = \sum_i dq_i \wedge dp_i$ relative to any set of local cotangent coordinates on M .¹⁴
3. If (M, g) and (M', g') are n -dimensional Riemannian manifolds, then for any $x \in M$ and $x' \in M'$, we know that g and g' endow the tangent spaces $T_x M$ and $T_{x'} M'$ with the same geometry; but in general we expect that no diffeomorphism $d : M \rightarrow M'$ will give an isometry between a neighbourhood of x and a neighbourhood of x' . But the Darboux theorem tells us that if (M, ω) is a finite-dimensional manifold equipped with a symplectic form,

¹²See section 3.2 below for further discussion and for an unpacking of this definition.

¹³Why regard p as a covector rather than a tangent vector? Because in general the momentum p of a system with Lagrangian L is defined as $\frac{\partial L}{\partial \dot{q}}$, which transforms as a covariant quantity under change of coordinates on Q .

¹⁴Where as above, a set of local coordinates on M induces a natural set of cotangent coordinates on T^*M . In example 7 below we will see a coordinate-free version of this construction that carries over to the infinite-dimensional case.

then (M, ω) is locally isomorphic to some $T^*\mathbb{R}^n$ equipped with its canonical cotangent bundle symplectic form. An immediate corollary is that every finite-dimensional symplectic manifold is even-dimensional.

Of course, for present purposes, the interest in identifying the symplectic structure lying behind the Hamiltonian version of the n -body problem lies in generalization. (1) Note that if we are interested in n particles interacting via forces that arise from a potential energy function V , as above, then we can construct a Hamiltonian treatment equivalent to the usual Newtonian one by again taking \mathcal{I} as the space of initial data, equipping it with the symplectic form ω as above, defining a Hamiltonian $H : \mathcal{I} \rightarrow \mathbb{R}$ as the sum of the kinetic and potential energies, and taking as our dynamical trajectories the integral curves of the vector field X_H on \mathcal{I} that solves $\omega(X_H, \cdot) = dH$. (2) More generally, we can model a vast number of classical mechanical systems as follows: let the space of initial states be a symplectic manifold (M, ω) (not necessarily a cotangent bundle) and let a Hamiltonian $H : M \rightarrow \mathbb{R}$ be given; then let the dynamics be given by the vector field X_H solving $\omega(X_H, \cdot) = dH$.

2.3 The Lagrangian Approach

It is helpful to approach the Lagrangian version of the n -body problem somewhat indirectly.¹⁵

Critical Points in Calculus

For $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the differential of f is given by $df = (\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_n})$. We say that f has a *critical point* at $x_0 \in \mathbb{R}^n$ if $df(x_0) = 0$; i.e., f has a critical point at x_0 if $df(x_0) \cdot e = 0$ for each $e \in \mathbb{R}^n$ (since \mathbb{R}^n is a linear space, we can identify $T_{x_0}\mathbb{R}^n$ with \mathbb{R}^n itself and let $e \in \mathbb{R}^n$ here). There are a number of helpful ways of thinking of $df(x_0) \cdot e$: (i) this quantity coincides with the *directional derivative* of f at x_0 in direction e ,

$$df(x_0) \cdot e = \lim_{t \rightarrow 0} \frac{f(x_0 + te) - f(x_0)}{t};$$

(ii) if we have a curve $\gamma : \mathbb{R} \rightarrow \mathbb{R}^n$ with $\gamma(0) = x_0$ and $\dot{\gamma}(0) = e$, then $df(x_0) \cdot e = \frac{d}{dt}f(\gamma(t))|_{t=0}$.

The Calculus of Variations and the Euler–Lagrange Equations

We now consider an infinite-dimensional analog: we look for critical points of a function defined on a space of curves in Euclidean space. This is the foundation of the Lagrangian approach to particle mechanics.

¹⁵For introductions to the Lagrangian approach via the calculus of variations, see [Dubrovin *et al.*, 1992, Chapter 6], [Lanczos, 1986, Chapters II and V], and [van Brunt, 2004]. For some of the rigorous underpinnings of the calculus of variations see, e.g., [Choquet-Bruhat *et al.*, 1977, §§II.A and II.B] and [Choquet-Bruhat and DeWitt-Morette, 1989, §II.3].

Let $Q = \mathbb{R}^n$, let $[a, b] \subset \mathbb{R}$ be a closed interval, and let $x, y \in Q$. Let $\Gamma(a, b; x, y)$ be the set of C^2 curves $\gamma : [a, b] \rightarrow Q$ with $\gamma(a) = x$ and $\gamma(b) = y$. And let $\Gamma(a, b; 0, 0)$ be the space of C^2 curves $\gamma : [a, b] \rightarrow Q = \mathbb{R}^n$ with $\gamma(a) = (0, \dots, 0)$ and $\gamma(b) = (0, \dots, 0)$. Both $\Gamma(a, b; x, y)$ and $\Gamma(a, b; 0, 0)$ are well-behaved infinite-dimensional spaces.¹⁶ For $\gamma \in \Gamma(a, b; x, y)$ we can think of $\Gamma(a, b; 0, 0)$ as $T_\gamma\Gamma(a, b; x, y)$ (think of $h \in \Gamma(a, b; 0, 0)$ as describing a vector field along γ).¹⁷

The tangent bundle of Q is $TQ = \mathbb{R}^{2n}$. Let $L : TQ \rightarrow \mathbb{R}$ be a smooth function. This allows us to define a function $I_{a,b} : \Gamma(a, b; x, y) \rightarrow \mathbb{R}$ by $I_{a,b}(\gamma) := \int_a^b L(\gamma(t), \dot{\gamma}(t)) dt$. We are interested in finding the critical points of $I_{a,b}$: these will be points in $\Gamma(a, b; x, y)$ (that is, curves $\gamma : [a, b] \rightarrow Q$) of special interest. Like any function on a well-behaved space, $I_{a,b}$ has a differential, which we denote $\delta I_{a,b}$; this can be thought of as one-form on $\Gamma(a, b; x, y)$.

DEFINITION 3 (Stationary Curves). We say that $\gamma : [a, b] \rightarrow Q$ is *stationary for L* over $[a, b]$ if $\delta I_{a,b}(\gamma) = 0$. We say that $\gamma : \mathbb{R} \rightarrow Q$ is *stationary for L* if its restriction to $[a, b]$ is stationary over $[a, b]$ for all closed intervals $[a, b]$.

As in the case of an ordinary function on \mathbb{R}^n , $\delta I_{a,b}(\gamma) = 0$ if and only if $\delta I_{a,b}(\gamma) \cdot h = 0$ for all $h \in T_\gamma\Gamma(a, b; x, y) = \Gamma(a, b; 0, 0)$. We can then calculate $\delta I_{a,b}(\gamma) \cdot h$ by finding $\frac{d}{d\varepsilon} I_{a,b}(\gamma[\varepsilon])|_{\varepsilon=0}$ for a curve $\gamma[\cdot] : \varepsilon \in \mathbb{R} \mapsto \gamma[\varepsilon] \in \Gamma(a, b; 0, 0)$ with $\gamma[0] = \gamma$ and $h = \frac{d}{d\varepsilon} \gamma[\varepsilon]|_{\varepsilon=0}$.

Let us calculate. Fix L and $[a, b]$. Let $\gamma \in \Gamma(a, b; x, y)$ and $h \in \Gamma(a, b; 0, 0)$. For each ε in some sufficiently small neighbourhood of zero, we define a curve $\gamma[\varepsilon] : \mathbb{R} \rightarrow Q$ by $\gamma[\varepsilon](t) := \gamma(t) + \varepsilon h(t)$. So $\gamma[\varepsilon]$ is a curve in $\Gamma(a, b; x, y)$ with $\gamma[0] = \gamma$ and with tangent $h = \frac{d}{d\varepsilon} \gamma[\varepsilon]|_{\varepsilon=0}$. Then:

$$\begin{aligned} \delta I_{a,b}(\gamma) \cdot h &= \frac{d}{d\varepsilon} I_{a,b}(\gamma[\varepsilon])|_{\varepsilon=0} \\ &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_a^b [L(\gamma[\varepsilon](t), \dot{\gamma}[\varepsilon](t)) - L(\gamma(t), \dot{\gamma}(t))] dt \\ &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \left(\int_a^b \varepsilon \left[\frac{\partial L}{\partial x}(\gamma(t), \dot{\gamma}(t)) \cdot h(\gamma(t)) \right. \right. \end{aligned}$$

¹⁶Let $\Gamma(a, b)$ be the space of C^2 curves $\gamma : [a, b] \rightarrow Q$. This is a linear space under pointwise addition (i.e., $(\gamma + \gamma')(x) = \gamma(x) + \gamma'(x)$) that can be made into a Banach space in a number of ways. $\Gamma(a, b; 0, 0)$ is a linear subspace of $\Gamma(a, b)$ while $\Gamma(a, b; x, y)$ is an affine subspace modelled on $\Gamma(a, b; 0, 0)$.

¹⁷We can think of $T_\gamma\Gamma(a, b; x, y)$ as being built as follows: one considers one-parameter family $\gamma[\varepsilon] : \varepsilon \in \mathbb{R} \mapsto \gamma[\varepsilon] \in \Gamma(a, b; x, y)$ of curves with $\gamma[0] = \gamma$, and declares such one-parameter families, $\gamma[\varepsilon]$ and $\gamma'[\varepsilon]$, to be equivalent if $\frac{d}{d\varepsilon} \gamma[\varepsilon]|_{\varepsilon=0} = \frac{d}{d\varepsilon} \gamma'[\varepsilon]|_{\varepsilon=0}$; $T_\gamma\Gamma(a, b; x, y)$ is the resulting space of equivalence classes. We construct a bijection between $\Gamma(a, b; 0, 0)$ and $T_\gamma\Gamma(a, b; x, y)$ thought of as the space of such equivalence classes by associating with $h \in \Gamma(a, b; 0, 0)$ the equivalence class containing $\gamma[\varepsilon] : \varepsilon \mapsto \gamma + \varepsilon \cdot h$.

$$\begin{aligned}
& + \left. \frac{\partial L}{\partial \dot{x}}(\gamma(t), \dot{\gamma}(t) \cdot \dot{h}(\gamma(t))) \right] + O(\varepsilon^2) dt \Big) \\
= & \int_a^b \frac{\partial L}{\partial x} \cdot h dt - \int_a^b \frac{d}{dt} \left(\frac{\partial L}{\partial \dot{x}} \right) \cdot h dt + \left(h \frac{\partial L}{\partial \dot{x}} \right) \Big|_a^b.
\end{aligned}$$

The first equality follows from a basic fact about the differential of a function; the second follows by definition; the third via Taylor's theorem; the fourth via an integration by parts. We now note that since h vanishes at $\gamma(a) = x$ and $\gamma(b) = y$ the third term in the final line vanishes. So

$$\delta I_{a,b}(\gamma_0) \cdot h = \int_a^b \left[\frac{\partial L}{\partial x} - \frac{d}{dt} \left(\frac{\partial L}{\partial \dot{x}} \right) \right] h dt.$$

To say that $\gamma : [a, b] \rightarrow Q$ is stationary over $[a, b]$ is to say that this expression vanishes for each h . So the condition that γ is stationary for L over $[a, b]$ is that the *Euler–Lagrange equation*

$$\frac{\partial L}{\partial x} - \frac{d}{dt} \frac{\partial L}{\partial \dot{x}} = 0 \tag{3}$$

holds along $\gamma(t)$ for $t \in [a, b]$. And the condition that $\gamma : [a, b] \rightarrow \mathbb{R}$ is stationary for L is just that the equation 3 is satisfied all along γ .

REMARK 4 (Parsing the Euler–Lagrange Equations). Here is how to unpack equation 3.¹⁸ Rewrite the expression for L , replacing \dot{x} everywhere by ξ . Then interpret equation 3 as a differential equation for admissible trajectories $x(t)$, understanding $\frac{\partial L}{\partial \dot{x}}$ to mean $\frac{\partial L(x, \xi)}{\partial \xi_i} \Big|_{\xi=\dot{x}(t)}$ and $\frac{d}{dt} \frac{\partial L}{\partial \dot{x}}$ to mean

$$\left(\frac{\partial^2 L}{\partial \xi_i \partial \xi_j} \ddot{x}_j + \frac{\partial^2 L}{\partial \xi_i \partial x_j} \dot{x}_j \right) \Big|_{\xi=\dot{x}(t)}.$$

Hamilton's Principle

Consider a physical system with configuration space Q (i.e., consider a system whose possible spatial positions are parameterized by Q). Let the kinetic energy be a function $T : TQ \rightarrow \mathbb{R}$ that arises from a Riemannian metric g on Q via $T(x, v) := g_x(v, v)$ and let $V : Q \rightarrow \mathbb{R}$ be a potential for a force depending on the configuration degrees of freedom alone. Then the Lagrangian for the system is $L(x, v) := T(x, v) - V(x)$. Hamilton's principle states that the stationary curves for L are the physically possible trajectories. Many physically interesting systems can be cast in this form — e.g., the n -body problem considered above. For such systems that Euler–Lagrange equations derived from Hamilton's principle are equivalent to the usual Newtonian equations of motion.

¹⁸For this, see, e.g., [Dubrovin *et al.*, 1992, 318].

Symplectic Structure of the Space of Solutions of the Euler–Lagrange Equations

Let Q be a manifold and TQ its tangent bundle. Let $L : TQ \rightarrow \mathbf{R}$ be a smooth function. The space, \mathcal{S} , of stationaries of L has a natural manifold structure: for those $\gamma \in \mathcal{S}$ defined at $t = 0$, we can take the values of $x = \gamma(0)$ and $v = \dot{\gamma}(0)$ relative to coordinates $\{x_i\}$ on Q as coordinates on \mathcal{S} ; doing this for each value of t gives us a differentiable atlas for \mathcal{S} . It follows that $\dim \mathcal{S} = \dim TQ$. We can also endow \mathcal{S} with a geometric structure: consider the boundary term, $h \frac{\partial L}{\partial \dot{x}}$, discarded above in deriving the Euler–Lagrange equations; since h is to be thought of as a tangent vector to the space $\Gamma(a, b; x, y)$, we must view $\alpha = \frac{\partial L}{\partial \dot{x}}$ as a one-form on that space; taking its exterior derivative gives us a two-form, $\omega := \delta\alpha$, on $\Gamma(a, b; x, y)$; the restriction of this form to \mathcal{S} is the structure we seek. In the coordinates we have introduced on \mathcal{S} , ω takes the form:

$$\omega = \frac{\partial^2 L}{\partial x^a \partial v^b} dx^a \wedge dx^b + \frac{\partial^2 L}{\partial v^a \partial v^b} dv^a \wedge dx^b.$$

For any L , this is a closed two-form. It is nondegenerate, and hence symplectic, so long as $\det \left[\frac{\partial^2 L}{\partial v^a \partial v^b} \right] \neq 0$.¹⁹ For Lagrangians of the form considered above this always holds — and we then find that (\mathcal{S}, ω) is (locally) symplectically isomorphic to the corresponding space of initial data that arises from a Hamiltonian treatment of the theory.²⁰

3 SYMPLECTIC MATTERS

Throughout the chapter, we are going to be investigating the representation of time and change in physical theories by asking about their representation in Lagrangian and Hamiltonian formulations of these theories. On the Lagrangian side, the focus is always on the space of solutions of the equations of our theory, while on the Hamiltonian side the focus is always on the space of initial data for those equations. It is a fact of primary importance that for well-behaved theories the space of initial data and the space of solutions share a common geometric structure — these spaces are isomorphic as symplectic manifolds. Thus the notion of a symplectic manifold and its generalizations will play a central role in our investigations.

It will be helpful to begin with a general discussion of the nature of symplectic manifolds: subsection 3.1 deals with some preliminary matters; subsection 3.2 offers a sketch of some of the basic concepts, constructions, and results of symplectic geometry as it figures in mechanics; subsection 3.3 offers the same sort of treatment of presymplectic geometry (a generalization of symplectic geometry that will play an important role in sections 6.2 and 7 below); subsection 3.4 discusses the sense in which a symplectic structure is the *sine non qua* of quantization.

¹⁹The symplectic structure of the space of solutions for Lagrangian theories is discussed in [Woodhouse, 1991, §§2.3 and 2.4].

²⁰This follows from the fact that Lagrangians arising from kinetic and potential terms of the sort considered above are always hyperregular; see, e.g., [Abraham and Marsden, 1978, 226].

3.1 Preliminaries

The spaces that we will come across below will be generalizations of ordinary n -dimensional manifolds in three respects. (i) They are allowed to be non-Hausdorff.²¹ (ii) They are allowed to be infinite-dimensional: a manifold is locally modelled on a vector space; we allow ours to be modelled on \mathbb{R}^n or on an infinite-dimensional Banach space.²² (iii) They are allowed to have mild singularities — roughly speaking, our spaces will be composed out of manifolds in the way that an ordinary cone is composed out of its apex (a zero-dimensional manifold) and mantle (a two-dimensional manifold) — but our spaces still have smooth structures and support tensors in much the same way that manifolds do.²³

In order to avoid becoming bogged down in technicalities, I will present my sketch of the required notions and constructions of symplectic and presymplectic geometry in the context of manifolds; but when in following sections I speak of ‘spaces’ rather than manifolds, it should be understood that I am allowing the spaces in question to have mild singularities of the sort mentioned above.

Below we will often be interested in the actions of Lie groups on manifolds, and in vector fields as the infinitesimal generators of such actions. Let me end this discussion of preliminary matters by reviewing some pertinent definitions and constructions.

Recall that a *Lie group* is a manifold which is also a group, with the operations of group multiplication, $(g, h) \in G \times G \mapsto g \cdot h \in G$, and the taking of inverses, $g \in G \mapsto g^{-1} \in G$, as smooth maps. An *action* of a Lie group G on a manifold M is a smooth map $\Phi : G \times M \rightarrow M$ such that: (i) $\Phi(e, x) = x$ for e the identity element of G and for all $x \in M$; (ii) $\Phi(g, \Phi(h, x)) = \Phi(gh, x)$ for all $g, h \in G$ and $x \in M$. One often writes $g \cdot x$ or $\Phi_g(x)$ for $\Phi(g, x)$.²⁴ The *orbit* through $x \in M$ of

²¹Recall that a topological space X is *Hausdorff* if for any distinct $x, y \in X$, there exist disjoint open sets U and V with $x \in U$ and $y \in V$. While most textbooks require manifolds to be Hausdorff, all of the basic constructions and results go through without this assumption — see [Lang, 1999]. As we will see in examples 32 and 33 below, the solution spaces of even the simplest physical systems can be non-Hausdorff.

²²[Abraham *et al.*, 1988] and [Lang, 1999] provide introductions to differential geometry that cover the case of infinite-dimensional Banach manifolds. See [Milnor, 1984, §§2–4] for an introduction to a more general approach, under which manifolds are modelled on locally convex topological vector spaces. Note that the inverse function theorem and the existence and uniqueness theorem for ordinary differential equations fail under this more general approach.

²³The spaces under consideration are Whitney stratified spaces. As suggested in the text, each such space admits a canonical decomposition into manifolds. This decomposition allows us to treat each point in such a space as lying in a manifold, which allows us to construct a space of tangent vectors and cotangent vectors at each point, and hence to construct tensors in the usual way. The dimensions of the manifold pieces (and of the tangent and cotangent spaces) will in general vary from point to point within the stratified space. See [Pflaum, 2001] or [Ortega and Ratiu, 2004, §§1.5–1.7] for a treatment of such spaces in the finite-dimensional case. The picture appears to be very similar in the infinite-dimensional examples that arise in physics: for general relativity, see [Andersson, 1989] and [Marsden, 1981, Lecture 10]; for Yang–Mills theories, see [Arms, 1981] and [Kondracki and Rogulski, 1986].

²⁴Equivalently, an action of G on M is a group homomorphism $g \mapsto \phi_g$ from G to $\mathcal{D}(M)$ (the group of diffeomorphisms from M to itself) such that the map $(g, x) \in G \times M \mapsto \phi_g(x) \in M$ is

the action is the set $[x] := \{g \cdot x : g \in G\}$. The action of a Lie group partitions a manifold into orbits.

While other Lie groups will figure below, we will most often be interested in the simplest of all Lie groups: the additive group \mathbb{R} . A *flow* on a manifold M is a one-parameter group of diffeomorphisms from M to M . So if $\{\Phi_t\}_{t \in \mathbb{R}}$ is a flow on M , then $\Phi_0(x) = x$ and $\Phi_t \circ \Phi_s(x) = \Phi_{t+s}(x)$ for all $x \in M$ and $s, t \in \mathbb{R}$. A flow $\{\Phi_t\}$ on M and an action $\Phi : \mathbb{R} \times M \rightarrow M$ of \mathbb{R} on M are more or less the same thing: given an \mathbb{R} -action $\Phi : \mathbb{R} \times M \rightarrow M$, one defines a flow $\{\Phi_t\}$ via $\Phi_t(x) = \Phi(t, x)$ for all $t \in \mathbb{R}$ and $x \in M$; likewise if one is given a flow and wants to define an \mathbb{R} -action.

Any \mathbb{R} -action on M induces a vector field X on M . Let $x \in M$ and consider the curve $\gamma_x(t) : \mathbb{R} \rightarrow M$ defined by $\gamma_x : t \mapsto \Phi_t(x)$. The image of γ_x in M is just the orbit $[x]$. Now suppose that $y \in [x]$ — i.e., there is $t \in \mathbb{R}$ such that $y = \Phi_t(x)$. Some facts follow immediately from the group property of $\{\Phi_t\}$. We find that the image of γ_y is also $[x]$ — so $[x] = [y]$. We find, in fact, that $\gamma_y(s) = \gamma_x(s+t)$ for all $s \in \mathbb{R}$; that is, each of the curves γ_y corresponding to points $y \in [x]$ agree up to choice of origin for their parameterization. So each orbit $[x]$ of our \mathbb{R} -action arises as the shared image of a (maximal) family of curves agreeing in their parameterization up to a choice of origin. As a convenient shorthand, we will speak of such a family of curves as an *affinely parameterized curve*, which we will think of as a curve with its parameterization fixed only up to a choice of origin. We can now construct a vector field X on M as follows: for $x \in M$ we define $X(x) = \dot{\gamma}_x(0)$ (the above discussion shows that X is a smooth vector field on M).

Now suppose that we are given a vector field X on a manifold M , and let us see whether we can think of X as generating an \mathbb{R} -action on M . Given $x \in M$, there is a unique curve γ_x passing through x at time $t = 0$ and such that for each value of t at which the curve is defined, its tangent vector at the point $\gamma(t) \in M$ is given by the value of X at that point.²⁵ Call this curve the *integral curve based at x* . We find that if y lies in the image of integral curve based at x , then the integral curves based at x and y have the same image and agree up to a choice of origin in their parameterization. So we might just as well replace these curves by the corresponding affinely parameterized curve, which we will call the *integral curve through x* (or y). So the vector field X allows us to define a family of integral curves on M , with each point in M lying on exactly one such curve. For $x \in M$ and $t \in \mathbb{R}$, let us agree that $\Phi(t, x)$ is the point that we reach by tracing t units along the integral curve through x , when this instruction is well-defined (recall that the integral curve based at x may only be defined on a subinterval of \mathbb{R}). This Φ will be an \mathbb{R} -action if and only if the domain of definition of each integral curve is all of \mathbb{R} . In this case, we call X a *complete* vector field, and call Φ the \mathbb{R} -action *generated by X* .

smooth.

²⁵This is just a statement of the existence and uniqueness theorem for first-order ordinary differential equations.

The picture is as follows: an \mathbb{R} -action Φ induces a vector field X on M , and X generates Φ . We think of the group $\{\Phi_t\}$ as consisting of the finite transformations generated by the infinitesimal transformations X (here and below, “infinitesimal” always means “living in the tangent space”). When X is incomplete, $\Phi(t, x)$ is not defined for all pairs (t, x) . In this case, Φ is known as a *local flow*. For many purposes, local flows are nearly as nice as flows, and it is still helpful to think of them as having vector fields as their infinitesimal generators.

3.2 Symplectic Manifolds

DEFINITION 5 (Symplectic Manifold). Let M be a manifold. A *symplectic form* on M is a closed nondegenerate two-form, ω . Here nondegeneracy means that at each $x \in M$ the map $\omega^b(x) : v \in T_x M \mapsto \omega(v, \cdot) \in T_x^* M$ is injective.²⁶ The pair (M, ω) is called a *symplectic manifold*.²⁷

DEFINITION 6 (Symplectic Symmetry). Let (M, ω) be a symplectic manifold. A *symplectic symmetry* of (M, ω) is a diffeomorphism $\Phi : M \rightarrow M$ that preserves ω in the sense that $\Phi^* \omega = \omega$ (i.e., the pullback of ω by Φ is just ω).

EXAMPLE 7 (Cotangent Bundle Symplectic Structure). Let Q be a finite- or infinite-dimensional manifold and let T^*Q be its cotangent bundle. We define a canonical symplectic form on T^*Q as follows. Let $\pi : T^*Q \rightarrow Q$ be the canonical projection $(q, p) \mapsto q$, and let $T\pi$ be the corresponding tangent map. There is a unique one-form θ on T^*Q such that $\theta(q, p) \cdot w = p(T\pi \cdot w)$ for all $(q, p) \in T^*Q$ and all $w \in T_{(q,p)} T^*Q$. We can then define the desired symplectic form as $\omega := -d\theta$, where d is the exterior derivative on T^*Q .²⁸

Let (M, ω) be a symplectic manifold, and let $C^\infty(M)$ be the set of smooth functions on M . For present purposes, the fundamental role of ω is to allow us to associate with each $f \in C^\infty(M)$ a smooth vector field X_f on M : X_f is implicitly defined by the equation $\omega(X_f, \cdot) = df$, where df is the exterior derivative of f (the nondegeneracy of ω guarantees that there is a unique solution to this equation).²⁹ We say that f *generates* X_f or that X_f is *generated by* f .

This basic construction has two fruits of the first importance:

²⁶Of course, for finite-dimensional M , $\omega^b(x)$ is surjective if and only if injective.

²⁷[Abraham and Marsden, 1978] and [Arnold, 1989] are the standard treatments of mechanics from the symplectic point of view. [Schmid, 1987] covers some of the same ground for the case of infinite-dimensional manifolds. [Ortega and Ratiu, 2004] is a comprehensive reference on the geometry and symmetries of finite-dimensional symplectic spaces (including singular spaces). [Cannas da Silva, 2006] is a helpful survey of symplectic geometry. [Weinstein, 1981] and [Gotay and Isenberg, 1992] offer overviews of the role of symplectic geometry in mathematics and physics.

²⁸In the finite-dimensional case and relative to a set of cotangent coordinates, ω is given by equation 2 above.

²⁹In the infinite-dimensional case, X_f may not be defined on all of M . For well-behaved f , we can deal with this by replacing M by the subspace on which X_f is defined. Below I will suppose that this has been done. For discussion, examples, and references see [Marsden, 1981, 11 ff.] and [Marsden and Ratiu, 1994, 106].

1. Via the map $f \mapsto X_f$, ω allows us to define a new algebraic operation on $C^\infty(M)$: the *Poisson bracket* between $f, g \in C^\infty(M)$ is $\{f, g\} := \omega(X_f, X_g)$.³⁰ This plays a crucial role in the theory of quantization — see section 3.4 below.
2. Via the map $f \mapsto X_f$, ω often allows us to associate smooth functions on M with one-parameter groups of symmetries of (M, ω) , and vice versa. (i) Let $f \in C^\infty(M)$ and let X_f be the vector field generated by f (via ω), and suppose that X_f is complete so that we are able to construct a corresponding flow, $\xi = \{\Phi_t\}_{t \in \mathbb{R}}$. Then each Φ_t preserves ω , in the sense that $\Phi_t^* \omega = \omega$.³¹ Furthermore: f itself is invariant under each Φ_t .³² (ii) Let $\xi = \{\Phi_t\}_{t \in \mathbb{R}}$ be a one-parameter group of symplectic symmetries of (M, ω) and let X be the vector field on M that is the infinitesimal generator of ξ . It is natural to ask whether we can find an $f \in C^\infty(M)$ that generates X . There are cases in which this is not possible.³³ But in the examples that arise in physics, this can typically be done. And by (i) above, when we can find such an f , we find that it is preserved by the flow ξ .³⁴

It is perhaps easier to grasp the function of a symplectic structure if one keeps in mind the Hamiltonian application of this framework.

DEFINITION 8 (Hamiltonian System). A *Hamiltonian system*, (M, ω, h) , consists of a symplectic manifold, (M, ω) , called the *phase space*, and a function $h : M \rightarrow \mathbb{R}$, called the *Hamiltonian*.

³⁰The Poisson bracket is a Lie bracket that obeys Leibniz's rule, $\{fg, h\} = f\{g, h\} + g\{f, h\}$.

³¹Indeed, we can further note that the Lie derivative of ω along X_f vanishes — and this holds even when X_f is incomplete. This provides a sense in which the local flow generated by an incomplete vector field preserves ω .

³²Indeed, the Lie derivative of f along X_f vanishes. This also holds when X_f is incomplete — so there is a sense in which the local flow generated by such an incomplete vector field preserves f .

³³See [Ortega and Ratiu, 2004, §4.5.16] for an example. See [Butterfield, this volume, §2.1.3] for further discussion.

³⁴More ambitiously, let G be a Lie group acting on M via symplectic symmetries, with $\dim G > 1$. Such a group will contain many one-parameter subgroups — as the group of isometries of Euclidean space contains a one-parameter group of translations corresponding to each direction in Euclidean space and a one-parameter group of rotations corresponding to each axis in Euclidean space. In this case, we can hope that for each one-parameter subgroup of G it is possible to find a function on M that generates that subgroup. If all goes very well — as it does in many examples that arise in physics — we can hope that the algebra of Poisson brackets between these generators will mirror the algebra of the group (i.e., there will be a Lie algebra isomorphism here). In this case, we speak of the existence of a *momentum map* (*warning*: terminology varies — many authors call these *infinitesimally equivariant momentum maps*). If f and g are functions on M such that their Poisson bracket vanishes, then we find that f is invariant under each symplectic symmetry in the one-parameter group generated by g . In particular, if G is a group of symplectic symmetries of (M, ω) and f a function on M such that the Poisson bracket of f with any function generating a one-parameter subgroup of G vanishes, then each of these generators is invariant under the one-parameter group of symmetries generated by f . [Woodhouse, 1991, §3.4] provides a useful guide to situations under which momentum maps are or are not available. See [Butterfield, this volume, §6] for further discussion.

We think of (M, ω) as the phase space of some physical system — such as the space of particle positions and momenta — and of h as assigning to each state of the system the total energy of that state. Together h and ω determine a flow $\{\Phi_t\}_{t \in \mathbb{R}}$ on M : each Φ_t maps each state to the state that dynamically follows from it after t units of time. h will be preserved under this group — this corresponds to the conservation of energy.³⁵

3.3 Presymplectic Manifolds

In sections 6.2 and 7 below we will be concerned with theories whose space of solutions and space of initial data are not symplectic.

DEFINITION 9 (Presymplectic Manifold). Let M be a manifold. A *presymplectic structure* on M is a closed degenerate two-form, ω ; we call (M, ω) a *presymplectic manifold*.³⁶ Here degeneracy means: at each point x there is a nontrivial *nullspace* $N_x \subset T_x M$ consisting of tangent vectors v such that $\omega_x(v, \cdot) = 0$.

A presymplectic structure ω on a manifold M induces a partition of M by submanifolds, $\{M_\alpha\}$ as follows. We define an equivalence relation on M by declaring $x, y \in M$ to be equivalent if they can be joined by a curve $\gamma : \mathbb{R} \rightarrow M$ each of whose tangent vectors is null — i.e., $\dot{\gamma}(t) \in N_{\gamma(t)}$ for each $t \in \mathbb{R}$. The equivalence classes, M_α , of this relation are called *gauge orbits*. For $x \in M$ we also denote the gauge orbit containing x by $[x]$. Each gauge orbit is a submanifold of M .³⁷ We call a function $f \in C^\infty(M)$ *gauge-invariant* if $f(x) = f(y)$ whenever x and y belong to the same gauge orbit of M (i.e., a function is gauge-invariant if and only if it is constant on gauge orbits).

We call a diffeomorphism from M to itself which preserves a presymplectic form ω a *presymplectic symmetry* of (M, ω) . We say that two presymplectic symmetries, Φ and Φ' , *agree up to gauge* if for each $x \in M$, $[\Phi(x)] = [\Phi'(x)]$ (i.e., for each $x \in M$, Φ and Φ' map x to the same gauge orbit); we call the set of presymplectic symmetries that agree with Φ up to gauge the *gauge equivalence class* of Φ . Similarly, we will say that two one-parameter groups, $\xi = \{\Phi_t\}$ and $\xi' = \{\Phi'_t\}$, of presymplectic symmetries agree up to gauge if Φ_t and Φ'_t agree up to gauge for each t ; the *gauge equivalence class* of $\xi = \{\Phi_t\}$ comprises all ξ' that agree with it up to gauge in this sense.

³⁵Often it will be possible to identify a larger group G of symplectic symmetries of (M, ω) that leaves h invariant (such as the group of Euclidean symmetries acting in the obvious way in Newtonian particle mechanics). Then a momentum map (see preceding footnote) would allow one to construct $\dim G$ independent quantities, whose algebra would mirror that of G , and that would be conserved under the dynamics generated by h .

³⁶Terminology varies: often (but not here) symplectic forms count as special cases of presymplectic forms; sometimes (but not here) presymplectic forms are required to have constant rank or to have well-behaved spaces of gauge orbits. On presymplectic geometry, see, e.g., [Gotay and Nester, 1980].

³⁷If X and Y are null vector fields on M (i.e., $X(x), Y(x) \in N_x$ for each $x \in M$) then, $[X, Y]$ is also a null vector field. It follows (by Frobenius' theorem) that the N_x form an integrable distribution, with the M_α as the integral manifolds.

If a presymplectic symmetry $\Phi : M \rightarrow M$ fixes each M_α (i.e., Φ maps points in M_α to points in M_α), then we call Φ a *gauge transformation*. Note that a gauge transformation agrees up to gauge with the identity map on M .

In the symplectic case: when all goes well, the equation $\omega(X_f, \cdot) = df$ allows one to associate each smooth function on a symplectic manifold (M, ω) with a one-parameter group of symplectic symmetries of (M, ω) — and vice versa.

In the presymplectic case: when all goes well, the equation $\omega(X_f, \cdot) = df$ allows one to associate each smooth gauge-invariant function on a presymplectic manifold (M, ω) with a gauge equivalence class of one-parameter groups of presymplectic symmetries of (M, ω) — and vice versa. So in the presymplectic case: if f generates the one-parameter group $\xi = \{\Phi_t\}$ of presymplectic symmetries via the equation $\omega(X_f, \cdot) = df$, then it also generates each $\xi' = \{\Phi'_t\}$ in the gauge equivalence class of ξ .

Note an interesting special case: any solution X_f of the equation $\omega(X_f, \cdot) = df$ for f a constant function is a vector field on M consisting of null vectors; so the corresponding one-dimensional group of presymplectic symmetries of (M, ω) consists of gauge transformations. Conversely: if $\xi = \{\Phi_t\}$ is a one-parameter group of gauge transformations of (M, ω) , then any function that generates ξ (via ω) is a constant function.

Given a presymplectic manifold (M, ω) , we can construct M' the space of gauge orbits of M . M' inherits a topological structure from M .³⁸ We will call the process of passing from M to M' *reduction*, and call M' the *reduced space*. In general, M' need not be a manifold, nor anything nearly so well-behaved as the spaces we want to consider below.³⁹ But when all goes well (as it usually does in the sort of cases considered below) M' will inherit from M a smooth structure (so it will be a space with at most mild singularities). And so long as some further technical conditions on ω hold, M' inherits from (M, ω) a two-form ω' that is nondegenerate as well as closed.⁴⁰ So, in this case, (M', ω') is a symplectic space. Note that each gauge-invariant $f \in C^\infty(M)$ corresponds to a unique $f' \in C^\infty(M')$. While f generates an equivalence class of one-parameter groups of presymplectic transformations of (M, ω) , f' generates a single one-parameter group of symplectic transformations of (M', ω') .⁴¹

³⁸We equip M' with the quotient topology, according to which a set $U' \subset M'$ is open if and only if $\pi^{-1}(U')$ is open in M (here π is the projection $x \in M \mapsto [x] \in M'$).

³⁹If (M, ω, H) is a Hamiltonian system in the sense of definition 8 above, then the restriction of ω to a surface, E , of constant energy is presymplectic — with the gauge orbits of $(E, \omega|_E)$ being the dynamical trajectories of the Hamiltonian system. If the dynamics is ergodic, then generic trajectories come arbitrarily close to each $x \in E$. It follows that the quotient space E' has the trivial topology, according to which the only open sets are the empty set and the space itself.

⁴⁰See [Marsden, 1981, 6] and [Ortega and Ratiu, 2004, §6.1.5].

⁴¹Each presymplectic symmetry of (M, ω) corresponds to a symplectic symmetry of (M', ω') , with two presymplectic symmetries corresponding to the same symplectic symmetry if and only if they agree up to gauge. Thus each gauge equivalence class of presymplectic symmetries corresponds to a single symplectic symmetry. And each gauge equivalence class of one-parameter groups of presymplectic symmetries corresponds to a single one-parameter group of symplectic

3.4 *Symplectic Structures and Quantization*

Quantization is the process of constructing a quantum counterpart to a given classical theory.⁴² As it is presently understood, it is a process which takes as its starting point a theory in Hamiltonian or Lagrangian form (or the discrete-time analog of such a theory). One does not know how to quantize a theory *qua* differential equations directly, without passing first to a Hamiltonian or Lagrangian recasting of the theory.⁴³

The following observations lend some plausibility to the idea that a symplectic structure is the *sine qua non* of quantization.

1. The core notion of quantization involves the following steps. One begins with a symplectic space (the space of classical solutions or initial data) and selects a set of functions on this space (classical physical quantities) that is closed under addition and the Poisson bracket induced by the symplectic structure. One then looks for a set of operators (quantum observables) acting on a space of quantum states, such that the algebra of these operators mirrors (or approximately mirrors, with increasingly better match as one approaches the classical limit) the algebra (under addition and the Poisson bracket) of the chosen classical quantities. One may then also need to take the further step of adding a Hamiltonian operator that implements the quantum dynamics.
2. Some classical theories have the unfortunate feature that when cast in Lagrangian or Hamiltonian form, they come to us with a space of solutions or initial data that is merely presymplectic. Typically, it is known that there is a symplectic space in the offing via reduction, as outlined above in section 3.3. But it may be difficult to construct this space, or it may happen, for one reason or another, that it seems easier to work with the presymplectic version of the theory. So a number of strategies have been developed for quantizing theories in presymplectic form: gauge fixing, Dirac constraint quantization, BRST quantization, etc. But it is very natural to think of each of these techniques as offering an indirect approach to the quantization of the underlying symplectic space.⁴⁴

However, there exist approaches to quantization that do not appear to employ the symmetries of the reduced space.

⁴²For overviews of the literature on quantization, see [Landsman, this volume] and [Ali and Englš, 2005].

⁴³For an attempt to show that in order for a theory to be quantizable, it must be derivable from a Lagrangian, see [Hojman and Shepley, 1991].

⁴⁴(i) Gauge fixing just amounts to finding a submanifold of the presymplectic space that is symplectically isomorphic to the reduced space; see, e.g., [Henneaux and Teitelboim, 1992, §1.4]. (ii) For Dirac's approach, see [Dirac, 2001] or [Henneaux and Teitelboim, 1992]. It is felt Dirac's algorithm should be amended when its output differs from the result of directly quantizing the reduced theory; see, e.g., [Duval *et al.*, 1990]. (iii) In the case of finite-dimensional systems, it can be shown that application of the BRST algorithm leads to a quantization of the reduced theory; see, e.g., [Loll, 1992] or [Tuyman, 1992]. (iv) On the relation between the BRST approach and the suggested amendment of the Dirac approach, see [Guillemin and Sternberg, 1990, §12].

symplectic structure of the classical spaces at all — for example Mackey quantization (which has a somewhat limited range of application) and path integral quantization (which has very wide application, but murky foundations in its application to field theories). As emphasized in [Landsman, this volume], the relation between the classical and the quantum is far from completely understood.

4 LAGRANGIAN FIELD THEORY

Differential equations are normally given to us in the following way. We are given a set of independent variables and a set of dependent variables, and a space of functions, \mathcal{K} , consisting of functions, u , that map values of the independent variables to values of the dependent variables. A differential equation Δ can be thought of as a condition on a function and its derivatives that is satisfied by only some $u \in \mathcal{K}$. We call the u that satisfy Δ the *solutions* to Δ and denote the space of such solutions by \mathcal{S} .

In physical applications, the independent variables typically parameterize space, time, or spacetime while the dependent variables parameterize the possible values of some quantity of interest. We can think of the functions $u \in \mathcal{K}$ as describing situations that are in some sense possible and of solutions $u \in \mathcal{S}$ as describing situations that are genuinely physically possible according to the theory whose laws are encoded in Δ . Although the terminology is not wholly perspicuous, I will speak of elements of \mathcal{K} as corresponding to *kinematical possibilities* and of elements of \mathcal{S} as corresponding to *dynamical possibilities*.

EXAMPLE 10 (Mechanics of a Particle). Consider the theory of a particle a particle in Euclidean space subject to a position-dependent force. The independent variable parameterizes time and the dependent variables parameterize the possible positions of the particle; an arbitrary continuous functions $x(t)$ of the form $t \in \mathbb{R} \mapsto x(t) \in \mathbb{R}^3$ describes a kinematically possible pattern of behaviour of the particle; $x(t)$ describes a dynamically possible behaviour if it satisfies the Newtonian equation $\ddot{x}(t) = F(x(t))$.

EXAMPLE 11 (The Klein–Gordon Field). The usual theory of a scalar field has the following ingredients: as independent variables we take inertial coordinates $\{t, x, y, z\}$ on Minkowski spacetime, V ; the theory has a single dependent variable, parameterizing the real numbers; so the kinematically possible fields are given by (suitably smooth) real-valued functions on Minkowski spacetime; the dynamically possible fields are those $\Phi : V \rightarrow \mathbb{R}$ satisfying the Klein–Gordon equation,

$$\frac{\partial^2 \Phi}{\partial t^2} - \frac{\partial^2 \Phi}{\partial x^2} - \frac{\partial^2 \Phi}{\partial y^2} - \frac{\partial^2 \Phi}{\partial z^2} - m^2 \Phi = 0.$$

Our primary concern below is with field theories — those physical theories whose laws are encoded in differential equations whose independent variables parameter-

ize spacetime.⁴⁵ We think of such a field theory as consisting of the following components: a spacetime V ; a space W in which the fields take their values, a space, \mathcal{K} , of kinematically possible fields (i.e., of functions from V to W satisfying suitable smoothness and boundary conditions); and a set of differential equations Δ .

This section has the following structure. In the first subsection below I discuss the Lagrangian approach, in which one singles out the set of dynamical possibilities within the space of kinematical possibilities via a variational problem for a Lagrangian rather than via the direct imposition of a differential equation. In the second subsection, I discuss a very important advantage of the Lagrangian approach over the direct approach: the former but not the latter allows one to equip the space of dynamical possibilities with a (pre)symplectic form. In the third subsection, I discuss the celebrated relation between conserved quantities and symmetries in the Lagrangian approach, first discerned by Noether. The discussion of these subsections is based upon [Zuckerman, 1987] and [Deligne and Freed, 1999, Chapters 1 and 2]; see also [Woodhouse, 1991, Chapters 2 and 7].

Before beginning it will be helpful to make some more specific assumptions about the theories we will be discussing. These assumptions will be in force throughout the remainder of the chapter.

Spacetime. Our spacetime V will always be an n -dimensional Hausdorff manifold V , with the topology $M \times \mathbb{R}$ for some $(n - 1)$ -manifold M . We will always think of time as having the topology of \mathbb{R} , so we will say that a spacetime with topology $M \times \mathbb{R}$ has *spatial topology* M . In particular, we will say that V is *spatially compact* if M is compact.

In most theories, the geometry of spacetime is fixed from solution to solution. So we typically think of V as carrying a solution-independent geometrical structure (I will be lazy, and sometimes use V to denote the manifold, sometimes the manifold and the geometry together).⁴⁶ Without worrying about precision, I will stipulate now that we will only be interested in spacetimes that are well-behaved. Examples: Newtonian spacetime, neoNewtonian spacetime, Minkowski spacetime, or other globally hyperbolic general relativistic spacetimes.

The spacetime geometries that we consider single out a distinguished class of hypersurfaces in V that correspond to instants of time.⁴⁷ In prerelativistic

⁴⁵*Prima facie*, the ontology of such a classical field theory satisfies Lewis's Humean supervenience — "the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing after another. . . . We have geometry: as system of external relations of spatiotemporal distance between points. . . . And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. . . . And that is all" [1986, ix f.]. Indeed, Lewis says that the picture was "inspired by classical physics" [1999, 226]. See [Butterfield, Unpublished] for doubts about the fit between Humean supervenience and classical physics.

⁴⁶We could also allow V to carry non-geometrical solution-independent structures, corresponding to external fields etc.

⁴⁷The crucial point is this: one needs to choose boundary conditions and a notion of instant

spacetimes the instants are just the hypersurfaces of absolute simultaneity. Typically, in relativistic spacetimes the instants are just the Cauchy surfaces. Occasionally in highly symmetric relativistic spacetimes, one requires instants to be Cauchy surfaces with nice symmetries — thus one might in some contexts require instants in Minkowski spacetime to arise as hypersurfaces of simultaneity relative to an inertial observer. Furthermore, in spacetimes carrying geometrical structures, it often makes sense to speak of certain curves in V as *possible worldlines of point-particles*. In prerelativistic spacetime, a curve counts as a possible worldline of a point-particle if it is transverse to the hypersurfaces of simultaneity; in relativistic spacetime such possible worldlines are given by timelike curves.

In addition to considering field theories set in fixed general relativistic background spacetimes, we also want to consider general relativity itself as a Lagrangian field theory. In that context the spacetime metric g is itself dynamical and varies from solution to solution. With this example in mind, we will allow a bare manifold of topology $M \times \mathbb{R}$ unequipped with any geometry to count as a spacetime for present purposes, so that general relativity can be developed alongside theories set in a fixed geometrical background. Note that even in a theory like general relativity in which the spatiotemporal geometry is dynamical we can still speak of a hypersurface as being an instant *relative* to a solution g .⁴⁸

Field Values. We will take W , the space of field values, to be a finite-dimensional vector space. However, we could afford to be more general, at the price of complicating some of the notation below. Our \mathcal{K} is a space of sections of a trivial vector bundle over V ; it follows that for $\Phi \in \mathcal{K}$ a tangent vector $\delta\Phi \in T_\Phi\mathcal{K}$ is also a map from V to W . We could have allowed \mathcal{K} to be a space of sections of an arbitrary fibre bundle $E \rightarrow V$. The chief complication that this would introduce is that a tangent vector $\delta\Phi \in T_\Phi\mathcal{K}$ would then be a section of the bundle $\Phi^*T(E/V)$.

Kinematically Possible Fields. In setting up a rigorous classical field theory, care must be taken in selecting differentiability and boundary conditions to impose on the kinematically possible fields. We can here afford to neglect such details, and just say that for each theory considered below, \mathcal{K} is taken to be a space of well-behaved functions $\Phi : V \rightarrow W$, required to satisfy appropriate conditions of differentiability and behaviour at infinity, but otherwise arbitrary. Note that while \mathcal{K} will be a manifold (often even an affine

in such a way that for certain $(n - 1)$ -forms, ω , for any instant $\Sigma \subset V$, $\int_\Sigma \omega$ converges, and is independent of the instant chosen (cf. fnn. 61 and 73 below). In the standard cases, the obvious notions of instant suffice.

⁴⁸Of course, in a theory in which the spacetime geometry is a solution-independent matter, $\Sigma \subset V$ is an instant relative to a solution Φ if and only if it is an instant according to the geometry of V .

or linear space), in general \mathcal{S} will be a nonlinear subspace of \mathcal{K} with mild singularities.

Differential Equations. The Lagrangian framework sketched below is very general and does not require a restriction on the order of the differential equations. However, because in later sections we will often be interested in comparing Hamiltonian and Lagrangian versions of the same theory, and because the Hamiltonian framework takes second-order equations as its point of departure, we will restrict attention to such equations beginning in section 5 below.

REMARK 12 (Finite-Dimensional Theories as Field Theories). In a classical theory of a system with finitely many degrees of freedom (finite systems of particles, rigid bodies, etc.) the configuration space Q is a manifold parameterizing the possible dispositions of the system in physical space. A history of the system is a curve $x : t \in \mathbb{R} \mapsto x(t) \in Q$. We can fit such theories into the present framework, by taking $W = Q$ and $V = \mathbb{R}$ (so the only independent variable is time). No harm comes of treating such a theory as a degenerate case of a field theory, so long as one does not forget that in this case the “spacetime” V parameterized by the independent variables of the theory is distinct from the spacetime in which the system is located.

REMARK 13 (Notation). Because a choice of V and W is implicit in a choice of \mathcal{K} , we can denote a field theory by (\mathcal{K}, Δ) .

4.1 The Lagrangian Approach

The role of the differential equations Δ of a theory is to cut down the space of kinematical possibilities \mathcal{K} to the space of dynamical possibilities \mathcal{S} .⁴⁹ The key insight of the Lagrangian approach is that for the vast majority of equations that arise in classical physics, there is an alternative way of singling out the subspace of solutions.⁵⁰

⁴⁹The text of this section is informal. More precise statements are given in the footnotes. The following terminology and results will be helpful.

The space $V \times \mathcal{K}$ is a manifold, and so carries differential forms and an exterior derivative operator. For $0 \leq p \leq n$ and $q \geq 0$ let $\Omega^{p,q}(V \times \mathcal{K})$ be the space of q -forms on \mathcal{K} that take their values in the space of p -forms on V : thus if $K \in \Omega^{p,q}(V \times \mathcal{K})$, $\Phi \in \mathcal{K}$, and $\delta\Phi_1, \dots, \delta\Phi_q \in T_\Phi\mathcal{K}$ then $K(\Phi, \delta\Phi_1, \dots, \delta\Phi_q)$ is a p -form on our spacetime V . Each differential form on $V \times \mathcal{K}$ belongs to some $\Omega^{p,q}(V \times \mathcal{K})$. Furthermore, we can write the exterior derivative, d , on $V \times \mathcal{K}$ as $d = D + \partial$, where D is the exterior derivative on V , mapping elements of $\Omega^{p,q}(V \times \mathcal{K})$ to elements of $\Omega^{p+1,q}(V \times \mathcal{K})$ (for $0 \leq p < n$), and ∂ is the exterior derivative on \mathcal{K} , mapping elements of $\Omega^{p,q}(V \times \mathcal{K})$ to elements of $\Omega^{p,q+1}(V \times \mathcal{K})$. We have $\partial D = -D\partial$.

Note that if $\Phi \in \mathcal{K}$ then a tangent vector $\delta\Phi \in T_\Phi\mathcal{K}$ is itself a map from V to W . So for each admissible p and q we can consider the subspace $\Omega_{loc}^{p,q}(V \times \mathcal{K}) \subset \Omega^{p,q}(V \times \mathcal{K})$ of local forms consisting of those K such that for any $\Phi \in \mathcal{K}$ and $\delta\Phi_1, \dots, \delta\Phi_q \in T_\Phi\mathcal{K}$, the value of the p -form $K(\Phi, \delta\Phi_1, \dots, \delta\Phi_q)$ at spacetime point $x \in V$ depends only on the values at x of $\Phi, \delta\Phi_1, \dots, \delta\Phi_q$, and finitely many of their derivatives.

⁵⁰For discussion of the scope of the Lagrangian approach, see [Bluman, 2005, §2.1].

DEFINITION 14 (Lagrangian). Let \mathcal{K} be a space of kinematically possible fields. A *Lagrangian*, L , on \mathcal{K} is a local map from \mathcal{K} to the space of n -forms on V (to say that L is *local* is to say that the value of $L(\Phi)$ at a point $x \in V$ depends only on the values at x of Φ and finitely many of its derivatives).⁵¹

Given a Lagrangian L , one can proceed, as in the treatment of the n -body problem sketched in section 2.3 above, to look for those kinematically possible Φ with the special property that infinitesimal perturbations at Φ make no difference to the value of $\int L(\Phi)$.

DEFINITION 15 (Variational Problem). Note that for each compact $U \subset V$, $S_U : \Phi \mapsto \int_U L(\Phi)$ is a real-valued function on \mathcal{K} . Let us call the assignment $U \mapsto S_U$ the *variational problem* of L .

DEFINITION 16 (Stationary Fields). We call $\Phi \in \mathcal{K}$ *stationary* for L if for each compact $U \subset V$ the effect of infinitesimally perturbing Φ inside U has no effect on the value of S_U .⁵²

DEFINITION 17 (Lagrangian Admitted by Δ). We call L a *Lagrangian for* (\mathcal{K}, Δ) if the set of Φ stationary for L coincides with the space \mathcal{S} of solutions of Δ . In this case we also say that Δ *admits* the Lagrangian L , and speak of \mathcal{S} as the space of solutions of (\mathcal{K}, L) .

REMARK 18 (Euler–Lagrange Equations). Given a Lagrangian, one can always find a set of equations Δ (the *Euler–Lagrange equations* for L) so that L is a Lagrangian for Δ . That is: a kinematically possible field $\Phi : V \rightarrow W$ is stationary for a Lagrangian L if and only if the Euler–Lagrange equations for L are satisfied. For Lagrangians depending only on the fields and their first-order derivatives, these equations require that

$$\frac{\partial L}{\partial \Phi^\alpha}(x_a) - \sum_{a=1}^n \frac{\partial}{\partial x_a} \left(\frac{\partial L}{\partial \Phi^\alpha_a} \right) (x_a) = 0 \tag{4}$$

hold at each point $x \in V$ (here a indexes coordinates on V , α indexes coordinates on W , and Φ^α_a stands for $\frac{\partial}{\partial x_a} \Phi^\alpha$).⁵³

⁵¹That is, $L \in \Omega_{loc}^{n,0}(V \times \mathcal{K})$.

⁵²That is, Φ is stationary for L if for each compact $U \subset V$ and for each $\delta\Phi \in T_\Phi \mathcal{K}$ whose support is contained in U we find that $\partial S_U(\delta\Phi) = \int_U \partial L(\Phi, \delta\Phi)$ vanishes. We can think of this as follows: fixing Φ , U , and $\delta\Phi$, we find a curve $\Phi[\varepsilon] : [-1, 1] \rightarrow \mathcal{K}$ such that $\Phi[0] = \Phi$ and $\frac{d}{d\varepsilon} \Phi[\varepsilon] |_{\varepsilon=0} = \delta\Phi$; the requirement that $\partial S_U(\delta\Phi) = 0$ amounts to $\frac{d}{d\varepsilon} \int_U L(\Phi[\varepsilon]) |_{\varepsilon=0} = 0$.

⁵³Of course, there is a coordinate-independent description of this. It is possible to show that $\partial L = E + DM$, where $E \in \Omega_{loc}^{n,1}(V \times \mathcal{K})$ and $M \in \Omega_{loc}^{n-1,1}(V \times \mathcal{K})$, with E determined uniquely by L and M determined up to the addition of an exact form DN , with $N \in \Omega_{loc}^{n-2,1}(V \times \mathcal{K})$. The condition that $\partial S_U(\delta\Phi) = 0$ becomes $\int_U E(\Phi, \delta\Phi) + DM(\Phi, \delta\Phi) = 0$ for all $\delta\Phi$ whose support is contained in U . Since $\delta\Phi$ vanishes along the boundary of U , Stokes’s theorem tells us that the second integrand makes no contribution. So Φ is stationary if and only if $\int_U E(\Phi, \delta\Phi) = 0$ for all such U and admissible $\delta\Phi$ — which is equivalent to saying that $E(\Phi, \delta\Phi) = 0$ for all such $\delta\Phi$. Relative to coordinates, this last equation is equivalent to equation 4 in the case of a Lagrangian depending only on first derivatives.

REMARK 19 (Trivially Differing Lagrangians). Let us say that Lagrangians L and L' differ trivially if L' is of the form $L'(\Phi) = L(\Phi) + \alpha_\Phi$ for each $\Phi \in \mathcal{K}$ with α_Φ an exact Φ -dependent n -form on V .⁵⁴ Let us say that if L and L' are Lagrangians, their variational problems $U \mapsto S_U$ and $U \mapsto S'_U$ are equivalent if for each compact $U \subset V$ and field $\Phi \in \mathcal{K}$, we have that any infinitesimal perturbation of Φ leaves the value of S_U unchanged if and only if it leaves the value of S'_U unchanged.⁵⁵ Lagrangians that differ trivially have equivalent variational problems.⁵⁶ It follows that trivially differing Lagrangians have the same space of solutions — indeed, they have the same Euler–Lagrange equations.⁵⁷

REMARK 20 (Uniqueness of Lagrangians). The previous remark shows that if Δ does admit a Lagrangian, it will admit infinitely many that differ trivially. Some Δ also admit multiple Lagrangians that do not differ trivially — e.g., the Newtonian equations for a particle moving in a spherical potential in three-dimensional Euclidean space.⁵⁸

REMARK 21 (Existence of Lagrangians). Not every set of equations Δ admits a Lagrangian.⁵⁹ A charged particle moving in the electromagnetic field of a magnetic monopole is an example of a system that does not admit a Lagrangian treatment.⁶⁰

⁵⁴I.e., $L' = L + DK$ where $K \in \Omega_{loc}^{n-1,0}(V \times \mathcal{K})$.

⁵⁵That is: the variational problems $U \mapsto S_U$ and $U \mapsto S'_U$ for Lagrangians L and L' are equivalent if for every compact $U \subset V$, every field $\Phi \in \mathcal{K}$, and every tangent vector $\delta\Phi \in T_\Phi\mathcal{K}$ with support contained in U , we have that $\partial S_U(\Phi)(\delta\Phi) = 0$ if and only if $\partial S'_U(\Phi)(\delta\Phi) = 0$.

⁵⁶Let $L' = L + DK$ with $K \in \Omega_{loc}^{n-1,0}(V \times \mathcal{K})$. Then for any compact $U \subset V$, $\Phi \in \mathcal{K}$, and $\delta\Phi \in T_\Phi\mathcal{K}$ with support contained in U , we have $\partial S_U(\Phi)(\delta\Phi) - \partial S'_U(\Phi)(\delta\Phi) = \int_U \partial DK(\Phi)(\delta\Phi)$. But $\partial D = -D\partial$, so the right hand side is $-\int_U D\partial K(\Phi)(\delta\Phi)$, which vanishes (by Stokes's theorem and the fact that $\delta\Phi$ vanishes on the boundary of U).

⁵⁷That is, if Lagrangians L and L' differ by a term of the form DK , then they share the same Euler–Lagrange operator E .

⁵⁸See [Crampin and Prince, 1988] and [Henneaux and Shepley, 1982] for this example. For field-theoretic examples, see [Nutku and Pavlov, 2002]. For a topological condition on $V \times W$ sufficient to ensure that Δ does not admit nontrivially differing Lagrangians, see [Anderson and Duchamp, 1980, Theorem 4.3.ii].

⁵⁹The problem of determining whether a given set of differential equations admits a Lagrangian is known as the *inverse problem of the calculus of variations* among mathematicians and as *Helmholtz's problem* among physicists. [Prince, 2000] is a helpful survey of results concerning finite-dimensional systems. [Anderson and Duchamp, 1980, §5] includes examples of field theories that do not admit Lagrangian formulations.

⁶⁰See [Anderson and Thompson, 1992, 4 f.]. For other examples, see [Prince, 2000].

4.2 The Structure of the Space of Solutions

The choice of a Lagrangian L allows us to equip \mathcal{S} with a closed two-form, Ω .⁶¹ So when Ω is nondegenerate, (\mathcal{S}, Ω) is a symplectic space; otherwise, it is presymplectic.⁶² Roughly speaking, one expects that Ω is nondegenerate if and only if the equations, Δ , of our theory have the property that specifying initial data determines a unique inextendible solution.⁶³

The choice of a Lagrangian brings into view the sort of structure required for the construction of a quantum theory. A set of differential equations Δ alone does not appear to determine such structure, and it is not known how to quantize a differential equation directly, without the introduction of a Lagrangian or a Hamiltonian. If Δ admits a Lagrangian L , then it also admits the whole class of Lagrangians that differ trivially from L (see remark 20 above). Unsurprisingly, trivially differing Lagrangians induces the same Ω on \mathcal{S} .⁶⁴ But when Δ admits Lagrangians L and L' that differ nontrivially, these Lagrangians can induce distinct geometric structures on \mathcal{S} ; and one expects that these distinct (pre)symplectic structures will lead to distinct quantizations of the given classical theory.⁶⁵ In the case mentioned above of a particle moving in a spherical potential, each of these elements is present: multiple nontrivially differing Lagrangians lead to distinct symplectic structures on the space of solutions, which lead in turn to physically distinct quantizations.⁶⁶

⁶¹Recall from footnote 53 above that we have the decomposition $\partial L = E + DM$, with E unique and M unique up to the addition of a D -exact form. We now define $Z := \partial M$. $Z \in \Omega_{loc}^{n-1,2}(V \times \mathcal{K})$ and is uniquely determined by L up to the addition of a term of the form DY , with $Y \in \Omega_{loc}^{n-2,2}(V \times \mathcal{K})$.

Let $\Phi \in \mathcal{S}$ be a solution, let $\delta\Phi_1, \delta\Phi_2 \in T_\Phi\mathcal{S}$, and let $\Sigma \subset V$ be an instant relative to Φ . Then we define $\Omega_\Sigma(\Phi, \delta\Phi_1, \delta\Phi_2) := \int_\Sigma Z(\Phi, \delta\Phi_1, \delta\Phi_2)$. We assume nice boundary conditions at infinity, so that Ω_Σ is well-defined, and so that replacing Z by $Z + DY$ makes no difference to Ω_Σ . We find that the value of Ω_Σ is independent of the instant chosen — because $Z(\Phi, \delta\Phi_1, \delta\Phi_2)$ is closed as an $(n-1)$ -form on V and we have been careful in our choice of notion of instant (see fn. 47). So we drop the subscript, and think of Ω as a two-form on \mathcal{S} , closed because Z is ∂ -exact.

⁶²Lagrange appears to have been the first to equip the space of solutions to a dynamical problem with this symplectic structure; see [Weinstein, 1981, §2], [Souriau, 1986], or [Iglesias, 1998].

⁶³As we will see below in 6.2, if the equations of motion admit gauge symmetries (so that uniqueness fails in a certain dramatic way), then Ω is presymplectic. I believe it is widely thought that this is the only way that Ω can fail to be symplectic — at least for the sort of examples that arise in physics.

⁶⁴Replacing L by $L + DK$ alters Z a term of the form DY , $Y \in \Omega_{loc}^{n-2,2}(V \times \mathcal{K})$. But because it is D -exact, this new term will not contribute to the integral over space that defines Ω (by Stokes's theorem and boundary conditions).

⁶⁵When (as in the Newtonian case) the equations of motion are second-order and the space of solutions is finite-dimensional, Lagrangians L and L' induce the same two-form on the space of solutions if and only if they differ trivially; see [Crampin and Prince, 1988, §II]. Presumably this in fact holds for a much wider range of cases.

⁶⁶See [Henneaux and Shepley, 1982].

4.3 Symmetries and Conserved Quantities

Given a set of equations Δ and a Lagrangian L admitted by Δ , there are three distinct notions of symmetry we might consider.⁶⁷ Roughly speaking, a *symmetry of Δ* is a map from \mathcal{K} to itself that fixes \mathcal{S} as a set and that is generated by an object local in the fields and their derivatives.⁶⁸ We can then consider the subset of *variational symmetries*, which also leave the variational problem of L invariant, or the subset of *Lagrangian symmetries* that leave L itself invariant. The three notions are distinct: every Lagrangian symmetry is a variational symmetry, but some theories have variational symmetries that are not Lagrangian symmetries; similarly, every variational symmetry is a symmetry of the associated equations of motion, but some equations that admit Lagrangians have symmetries that are not variational symmetries of any Lagrangian for the theory.⁶⁹

For present purposes, it is natural to focus on variational symmetries of physical theories. For, on the one hand, the class of Lagrangian symmetries excludes some physically important symmetries — and in any case, within the Lagrangian approach it is not clear that it is more natural to focus on symmetries of the Lagrangian than on symmetries of the variational problem. On the other hand, the class of symmetries of equations that are not variational symmetries does not appear to include any symmetries of absolutely central physical interest — and it is at the level of variational symmetries (rather than symmetries of equations) that the powerful results of Noether, cementing a connection between certain special types of one-parameter groups of variational symmetries and certain special types of conserved quantities in classical field theories, are naturally situated.⁷⁰

Here is a statement of the central result. Let us call a one-parameter group, $\xi = \{g_t\}$, of diffeomorphisms from \mathcal{K} to itself a *Noether group for L* if its infinitesimal generator leaves invariant the variational problem of L and is local in the appropriate sense.⁷¹ Given a Noether group $\xi = \{g_t\}$ for (\mathcal{K}, L) , there is

⁶⁷See [Olver, 1993, Chapters 2, 4, and 5] for the relevant notions. *Warning:* terminology varies — sometimes my *Lagrangian symmetries* are called *variational symmetries*, sometimes my *variational symmetries* are called *divergence symmetries*, etc.

⁶⁸See [Olver, 1993, §5.1] for details.

⁶⁹The wave equation in (2+1) dimensions has a dilational symmetry that is not a variational symmetry and inversion symmetries that are variational but not Lagrangian; see [Olver, 1993, Examples 2.43, 4.15, 4.36, and 5.63]. Example 4.35 of the same work shows that Galilean boosts are variational symmetries for the n -body problem but are not Lagrangian symmetries. Indeed, no Lagrangian for Newtonian particles subject to forces derived from a potential can be invariant under the full group of symmetries of neoNewtonian spacetime; see [Souriau, 1997, Remark 12.136].

⁷⁰Note, however, that there do exist results establishing links between symmetries of equations with conserved quantities, without detouring through the Lagrangian framework; see, e.g., [Bluman, 2005].

⁷¹More precisely, let ξ be a one-parameter group of diffeomorphisms from \mathcal{K} to itself and let X be the corresponding vector field on \mathcal{K} (i.e., X is the vector field whose flow is ξ). ξ is a Noether group if the following two conditions hold. (i) X is an *infinitesimal variational symmetry of L* : there exists an $R \in \Omega_{loc}^{n-1,0}(V \times \mathcal{K})$ such that $\partial L(\Phi, (X(\Phi))) = DR(\Phi)$ for all $\Phi \in \mathcal{S}$. (ii) X is *local*: for any $\Phi \in \mathcal{K}$, $X(\Phi) \in T_{\Phi}\mathcal{K}$ is local on V , in the sense that at any point $x \in V$, we find that $X(\Phi)(x)$ depends only on the value at x of Φ and finitely many of its derivatives (recall that

a map J_ξ , called the *Noether current* associated with ξ , that maps solutions to $(n-1)$ -forms on V .⁷² Given an arbitrary solution, $\Phi \in \mathcal{S}$, and an instant $\Sigma \subset V$ we integrate $J_\xi(\Phi)$ over Σ to give the *Noether charge*, $Q_{\xi,\Sigma}(\Phi) := \int_\Sigma J_\xi(\Phi)$. We note that $Q_{\xi,\Sigma}(\Phi)$ is independent of the Σ chosen (so long as the integral is well-defined!).⁷³ That is: $Q_{\xi,\Sigma}(\Phi)$ is a quantity that is constant in time within the solution (V, Φ) . Thus we might as well denote it simply $Q_\xi(\Phi)$, and think of the Noether charge, Q_ξ , associated with ξ as a function on \mathcal{S} .

REMARK 22 (Noether Charges Generate Symmetries). Since Ω is a closed two-form, (\mathcal{S}, Ω) is a symplectic or presymplectic space: so the results of sections 3.2 and 3.3 above apply. As one would expect, Q_ξ is in fact the symplectic/a presymplectic generator of the one-parameter group ξ (thought of now as acting on \mathcal{S}). The beauty of Noether's result is that it shows how to construct the generator of ξ via the integration of local objects on spacetime.

REMARK 23 (Trivial Conservation Laws). So far, nothing we have said guarantees that Q_ξ is an *interesting* function on \mathcal{S} — it might, for instance be a the zero function, if $J_\xi(\Phi)$ is exact as an $(n-1)$ -form on V . Such trivial Noether charges do in fact occur when Ω is presymplectic and ξ is a group of gauge transformations. We will see examples of this in section 6.2 below.

5 TIME AND CHANGE IN WELL-BEHAVED FIELD THEORIES

Turn we now to the representation of time and change in physical theories. In the remaining sections Hamiltonian formulations of theories will play an important role. So we henceforth restrict attention to theories with second-order equations of motion.

In this section, we discuss ideally well-behaved theories. We impose three further assumptions, which are in effect jointly for this section only: (a) global existence of solutions; (b) uniqueness of solutions; (c) our spacetime admits a time translation symmetry under which the variational problem of our Lagrangian is invariant.

We will see that in this context, we have three \mathbb{R} -actions: a notion of time translation on spacetime; a notion of time translation on the space of solutions of the theory; and a notion of time evolution on the space of initial data of the theory. We also find that the space of solutions and the space of initial data are isomorphic as symplectic spaces, and that there is a natural intertwining of the notion of time

an element of $T_{\mathbb{F}}\mathcal{K}$ is itself a map from V to W).

⁷²The Noether current associated to L and ξ is the element of $J_\xi \in \Omega_{loc}^{n-1,0}(V \times \mathcal{K})$ given by $J_\xi(\Phi) := R(\Phi) - M(\Phi, X(\Phi))$, where X is the infinitesimal generator of ξ , R is the object introduced in the preceding footnote, and M is the object introduced in footnote 53.

⁷³ Because $J_\xi(\Phi)$ is closed as an $(n-1)$ -form on V and because we have been careful in our choice of notion of instant (see fn. 47). Note, in fact, that so long as $\Sigma, \Sigma' \subset V$ are compact $(n-1)$ -manifolds that determine the same homology class in V , we will have $\int_\Sigma J_\xi(\Phi) = \int_{\Sigma'} J_\xi(\Phi)$ (see, e.g., [Lee, 2003, 431] and [Lee, 2000, 300 f.] for relevant notions and results). Hence we get a sort of conservation law even if, e.g., Σ and Σ' are not spacelike according to the geometry of V . See [Torre, Unpublished] for an introduction to such conservation laws.

translation on the space of solutions with the notion to time evolution on the space of initial data. So in this domain one can say simply (if awkwardly) that time is represented as a symmetry of the laws — and leave it open whether one means time translation or time evolution, since in the end the two come to much the same thing.

This section has five subsections. The first is devoted to the Lagrangian picture, the second to the Hamiltonian, the third to the relation between these pictures, the fourth to a discussion of the representation of time and change. The final subsection offers an overview.

5.1 The Lagrangian Picture

Let us be more precise about the special assumptions in play in this section. We impose the following conditions on our spacetime V , equations of motion Δ , and Lagrangian L .

Global Existence of Solutions. We assume that each admissible set of initial data for Δ is consistent with a solution defined on all of V .⁷⁴

Uniqueness of Solutions. If Φ and Φ' are solutions that agree in the initial data that they induce on an instant $\Sigma \subset V$, then they agree at any point $x \in V$ at which they are both defined.

Time Translation Invariance of the Lagrangian. We require our spacetime V to have a nontrivial geometrical structure, strong enough to single out a class of $(n - 1)$ -dimensional submanifolds that count as instants and a class of one-dimensional submanifolds that count as possible worldlines of point-particles. Let $\bar{\xi} = \{\bar{g}_t\}$ be a one-parameter group of spacetime symmetries of V , and consider the orbits of $\bar{\xi}$ in V (the orbit $[x]$ of $\bar{\xi}$ through $x \in V$ is the image of the curve $x(t) := \bar{g}_t \cdot x$). We call $\bar{\xi}$ a *time translation group for V* if the orbits of $\bar{\xi}$ are possible worldlines of point-particles according to the geometry of V ; in this case, we call these orbits *worldlines adapted to $\bar{\xi}$* . We will typically denote time translation groups as $\bar{\tau}$.

Let \bar{G} be a group of spacetime symmetries of V . Given $\bar{g} \in \bar{G}$ we can define a diffeomorphism $g : \mathcal{K} \rightarrow \mathcal{K}$ via $g(\Phi(x)) = \Phi(\bar{g}^{-1} \cdot x)$. In decent Lagrangian theories, one expects that if $\bar{\xi} = \{\bar{g}_t\}$ is a one-parameter group of spacetime symmetries, then $\xi = \{g_t\}$ is a Noether group for L . In this situation, ξ will map solutions to solutions; so that each $g_t \in \xi$ restricts to a map from \mathcal{S} to itself; these maps are symplectic automorphisms of (\mathcal{S}, Ω) (I won't bother introducing notation to distinguish between the action of ξ on \mathcal{K} and the restriction of this action to \mathcal{S}). In this section we assume that each time

⁷⁴Since we are restricting attention to theories with second-order equations of motion, specifying initial data involves specifying the field values and their time-rate of change at some initial instant.

translation group, $\bar{\tau}$, of V gives rise in this way to a Noether group, τ , of L . I will call such a τ a *dynamical time translation group*.

Within the class of theories that arise in physics, it appears to be an immediate consequence of the uniqueness assumption that the form Ω induced by L on the space of solutions is nondegenerate, and hence symplectic. We denote by H the corresponding conserved quantity guaranteed by Noether's theorem (in physically realistic theories, H arises by integrating the stress-energy of the field over an arbitrary instant).⁷⁵ Of course, H generates, via Ω , the action of τ on \mathcal{S} .

EXAMPLE 24 (Field Theory in Newtonian Spacetime). In Newtonian spacetime, each symmetry can be written as the product of a time translation with an isometry of absolute space. In coordinates adapted to the privileged absolute frame, we can write points of spacetime as (t, \mathbf{x}) . Then the (orientation-preserving) symmetries of V are of the form $(t, \mathbf{x}) \mapsto (t + s, R(\mathbf{x}) + \mathbf{c})$, where $s \in \mathbb{R}$ implements a time translation, R is a matrix implementing a rotation in absolute space, and $\mathbf{c} \in \mathbf{R}^3$ implements a spatial translation. Up to a choice of temporal unit, there is a unique time translation group, $\bar{\tau} : (t, \mathbf{x}) \mapsto (t + s, \mathbf{x})$; the worldlines of the points of absolute space are adapted to this group. We are supposing that the corresponding group τ acting on the space of solutions is a dynamical time translation group. The Noether charge associated with τ , $H : \mathcal{S} \rightarrow \mathbb{R}$, assigns to each solution the total energy of the system at any instant (since we are considering a theory invariant under time translations, the value of the total energy along a slice is a constant).⁷⁶

EXAMPLE 25 (Field Theory in Minkowski Spacetime). The symmetry group of Minkowski spacetime is the Poincaré group. Each inertial frame picks out a notion of simultaneity, and a time translation group, $\bar{\tau}$; the worldlines of observers at rest in the chosen frame will be adapted to this group. (Equivalently, such group is determined by the choice of a timelike vector in spacetime.) In Poincaré-invariant field theories we can choose inertial coordinates (t, x_1, x_2, x_3) such that our chosen $\bar{\tau}$ acts via $(t, x_1, x_2, x_3) \mapsto (t + s, x_1, x_2, x_3)$. In such coordinates, the Noether current is just the component T^{00} of the stress-energy tensor of the field — the Noether charge being given, as always, by the integral of the Noether current over any instant.⁷⁷

⁷⁵For the stress-energy tensor and its role in the examples below, see [Choquet-Bruhat and DeWitt-Morrette, 1989, §II.7] and [Deligne and Freed, 1999, §2.9].

⁷⁶The Noether charge generating spatial translation (rotation) in a given direction (about a given axis) assigns to a solution the corresponding component of the linear (angular) momentum of the system at an instant. In fact, we get a momentum map (see fn. 34) for the action of the group of symmetries of Newtonian spacetime — the Poisson bracket algebra of the Noether charges mirrors the Lie bracket relations between the infinitesimal generators of the corresponding one-parameter groups. It is impossible, however, to construct a momentum map for the symmetry group of neo-Newtonian spacetime; see [Woodhouse, 1991, §3.4] for this and other examples in which the construction of a momentum map is impossible.

⁷⁷One can again construct a momentum map (see fn. 34) — with spacelike translations generated by the components of linear momentum, etc., in the familiar way.

EXAMPLE 26 (Field Theories in a Curved Spacetime). While a generic general relativistic spacetime admits no non-trivial symmetries, a solution in which, intuitively, the geometry of space is constant in time admits a time translation group. Let V be a globally hyperbolic and time-oriented general relativistic spacetime that possess such a $\bar{\tau}$. Let X_a be the vector field tangent to the orbits of $\bar{\tau}$ (so X is a timelike Killing field). Let $T^{ab}(\Phi)$ be the stress-energy tensor of the field Φ and suppose that $\nabla_a T^{ab} = 0$ (this typically holds in cases of physical interest). Let $\Sigma \subset V$ be an instant (i.e., a Cauchy surface) and let n_a be the field of unit future-pointing normal vectors along Σ . We can define the energy-momentum vector of T^{ab} relative to X_a as $P^b := X_a T^{ab}$ and define the energy along Σ as $\int_{\Sigma} P^a n_a dx$. This last quantity is in fact the Noether charge, and is independent of Σ .

5.2 The Hamiltonian Picture

The basic idea behind the Hamiltonian approach is to work with the space of initial data of the equations of the theory rather than with the space of solutions to the equations — roughly and heuristically speaking, this means working with the space of instantaneous states of the theory rather than with its space of possible worlds.

Deterministic equations of motion tell us what the state of the system must be at earlier and later times if it is in a given initial state. So, at least for well-behaved equations of motion, the dynamical content of the equations of motion ought to be encodable in a flow on the space of initial data, with the integral curves of this flow being the dynamically possible trajectories through the space of instantaneous states.

The special assumptions in play in this section (global existence and uniqueness of solutions and the presence of a dynamical time translation group) imply (at least for the sort of the theories that arise in physics) that the space of initial data carries a symplectic structure that generates the dynamics of the theory when supplemented by the function that assigns to an initial data set the total energy of a system in that state. The dynamics can be thought of as encoded in an \mathbb{R} -action on the space of initial data that implements time evolution. As we will see, these structures on the space of initial data — symplectic structure, Hamiltonian, and group action — are all closely related to the corresponding objects on the space of solutions that arise on the Lagrangian side.

Intuitively speaking, an instantaneous state of the field is a specification at each point of space of the value of the field and its time rate of change; and in giving a sequence of such instantaneous states, we describe how the values of these variables evolve through time at each point of space. So in order to construct a Hamiltonian formulation of a theory in which the total history of a system is described via a trajectory through the space of initial data, we need to effect some sort of notional decomposition of spacetime into space and time.⁷⁸

⁷⁸Note that we did not require any such decomposition in setting up the Lagrangian formalism in section 4 above. It is, of course, crucial to distinguish the symplectic space of solutions that arises in this formalism from the symplectic velocity phase space that arises in some Lagrangian

Informally, we can picture what is required as being a preferred family of observers together with a notion of simultaneity. Spacetime is partitioned by the worldlines of these observers (these need not be at rest relative to one another, but we do require that the worldlines involved be possible worldlines of point-particles according to the geometry of V). Each observer carries a clock; and we suppose that the set of points at which these clocks read $t = 0$ forms an instant in V . We call such a set of observers equipped with such a notion of simultaneity a *frame*. We say that a frame is *adapted* to the time translation group $\bar{\tau}$ when the following two conditions obtain. (i) The worldline of each observer is an orbit of the group $\bar{\tau}$ acting on V . (ii) Up to a choice of origin and a choice of unit of measurement, $\bar{\tau}$ gives us a parameterization of the set of instants of V , which allows us to determine ratios of temporal intervals; we require that the clock readings respect these determinations.

Still speaking informally, we can say that relative to a choice of frame, the state of the field at time t is an assignment of field value and momentum to each observer (i.e., the values of the field and its time rate of change at the spacetime point the observer occupies at the given instant), and that a history of the field is an assignment of these quantities along the worldline of each observer. So we can take an initial data set to be a pair of functions (corresponding to the field value and its time rate of change) defined on the space of observers of our frame — this space acts as a sort of abstract instant, with the same topological and geometrical structure as the concrete instants that arise as subsets of V .

We can make this more precise as follows.

DEFINITION 27 (Slicing). Let V be a spacetime with geometry and let S be an $(n-1)$ -dimensional manifold (possibly carrying a Riemannian metric). Then an *S -slicing* of V is a diffeomorphism $\sigma : \mathbb{R} \times S \rightarrow V$ such that: (i) each $\Sigma_t := \sigma(\{t\} \times S)$, $t \in \mathbb{R}$, is an instant in V (with σ providing an isomorphism between the geometry of Σ_t and the geometry of S , if any); (ii) each $X_x := \sigma(\mathbb{R} \times \{x\})$, $x \in S$, is a possible worldline of a point-particle according to the geometry of V . We call S the *abstract instant* of σ and each Σ_t an *instant in the slicing*. When V admits a time translation group $\bar{\tau}$, we call a slicing σ of V *adapted* to $\bar{\tau}$ if the following conditions are met: (a) each X_x is an orbit of $\bar{\tau}$; (b) any two instants of the slicing are related by a time translation in $\bar{\tau}$; (c) up to a choice unit and origin, the parameterizations of each X_x given by σ and by $\bar{\tau}$ agree.

EXAMPLE 28 (Newtonian Slicings). In Newtonian spacetime there is of course a unique partition of spacetime by instants and (up to a choice of unit) a unique time translation group $\bar{\tau}$. Furthermore, in this setting it is possible to take S to be the space of worldlines of the points of absolute space.⁷⁹ So the only freedom in constructing a slicing adapted to $\bar{\tau}$ is in choosing an origin and a unit for the

treatments — the latter does, while the former does not, presuppose a division of spacetime into instants. Cf fn. 6 above.

⁷⁹This space comes equipped with a natural Euclidean structure — since the distance between points of absolute space is constant in time, we can define the distance between two worldlines of such points to be the distance between the points.

parameterization of the instants by the reals.

EXAMPLE 29 (Flat Minkowski Slicings). In the setting of Minkowski spacetime it is sometimes natural to restrict attention to instants which arise as surfaces of simultaneity for inertial observers. In this case, our abstract instant S will again have the structure of Euclidean space. In order to construct a slicing, we must choose an instant $\Sigma_0 \subset V$ corresponding to $t = 0$, an isometry from S to Σ_0 , a unit of temporal measurement, and a notion of time translation associated to an inertial observer.

EXAMPLE 30 (Generic Minkowski Slicings). More generally, in the Minkowski spacetime setting it is possible to allow arbitrary Cauchy surfaces to count as instants — in this case, one will choose S to have some non-trivial Riemannian geometry. Now there is a truly vast — indeed, infinite-dimensional — family of instants to choose from (as we allow the geometry of S to vary). On the bright side, a generic instant admits no nontrivial isometries — so having chosen S and Σ there will be no freedom in constructing an isometry from one to the other.

Let us consider a Lagrangian theory satisfying all of the present conditions, and fix a slicing of V adapted to a notion of time translation, $\bar{\tau}$, that gives rise to a dynamical time translation group τ . We can then construct a Hamiltonian version in the following steps.

1. *Given an instant and a solution, construct the instantaneous field configuration and momentum.* Let Σ be an instant contained in the given slicing and let $\Phi : V \rightarrow W$ be a solution. We define $\phi : \Sigma \rightarrow W$, the field configuration on Σ , by $\phi := \Phi|_{\Sigma}$. And we define $\dot{\phi} : \Sigma \rightarrow W$, the field velocity on Σ , as follows: at each $x \in \Sigma$, $\dot{\phi}(x)$ is the rate of change at x of the field values along the orbit of $\bar{\tau}$ through x .⁸⁰ In order to construct the instantaneous momentum of the field, we apply the usual recipe for constructing canonical momentum variables, defining $\pi := \frac{\partial L}{\partial \dot{\phi}}$ (π is a map from Σ to W^* , the vector space dual to W).
2. *Given the instantaneous field configuration and momentum, construct the corresponding initial data.* This is just a matter of using σ to pull back ϕ and π , so that we can regard them as functions on S rather than Σ . Sloppily, I will use the same names for initial data defined on S and the corresponding functions defined on $\Sigma \subset V$.
3. *Construct the space of initial data, \mathcal{I} .* Let \mathcal{Q} be the space of all $\phi : S \rightarrow W$ that can arise via the previous two steps as we allow Φ to vary in \mathcal{S} .⁸¹ The set of all pairs (ϕ, π) that can arise via these steps is just cotangent bundle, $T^*\mathcal{Q}$. This space is the space, \mathcal{I} , of initial data for our theory. It carries a canonical symplectic structure, ω (see example 7 above).

⁸⁰That is, let $x_0 \in \Sigma$ and find $y_0 \in S$, $t_0 \in \mathbb{R}$ such that $\sigma(t_0, y_0) = x_0$ and define the curve $x : \mathbb{R} \rightarrow V$ by $x(t) := \sigma(t, y_0)$; then let $\dot{\phi}(x) := \lim_{h \rightarrow 0} \frac{1}{h} (\phi(x(t_0 + h)) - \phi(x_0))$.

⁸¹Allowing Σ and the slicing to vary as well would make no difference in the present case, so long as S and its geometry are held fixed.

4. *Construct a Hamiltonian.* We define $h : \mathcal{I} \rightarrow \mathbb{R}$, the Hamiltonian on the space of initial data, as follows. Let $(\phi, \pi) \in \mathcal{I}$ be initial data and let $\Sigma \subset V$ be an instant (not necessarily one in our slicing). Let Φ be the solution that induces (ϕ, π) on Σ and define $h(\phi, \pi) := \int_{\Sigma} \pi(x)\dot{\phi}(x) - L(\Phi)(x)dx$ (in the present context, the result does not depend on the instant Σ chosen).⁸²
5. *Construct the Dynamics.* Together h and ω determine a vector field χ on \mathcal{I} that encodes the dynamics of our theory. The integral curves of χ are the possible dynamical trajectories — if the state is (ϕ_0, π_0) at time $t = 0$, then the state t units of time later can be found by tracing t units of time along the integral curve passing through (ϕ_0, π_0) . This gives us a flow on \mathcal{I} , which preserves both ω and h (the flow is global rather than local because we are assuming that solutions are defined for all values of t).

5.3 Relation between the Lagrangian and Hamiltonian Pictures

For each instant $\Sigma_t := \sigma(\{t\} \times S)$ in our slicing σ , we define $T_{\Sigma_t} : \mathcal{S} \rightarrow \mathcal{I}$ to be the map that sends a solution Φ to the initial data set $(\phi, \pi) \in \mathcal{I}$ that results when the slicing σ is used to pullback to S the initial data induced by Φ on Σ_t . Because we are assuming global existence and uniqueness for solutions given initial data, T_{Σ_t} is in fact a bijection. Indeed, it is a diffeomorphism. Furthermore, $T_{\Sigma_t}^* \omega = \Omega$, so each T_{Σ_t} is in fact a symplectic isomorphism between (\mathcal{S}, Ω) and (\mathcal{I}, ω) .

Note that in typical theories distinct instants in the slicing lead to distinct isomorphisms. If Σ_t and $\Sigma_{t'}$ are instants in our slicing and $T_{\Sigma_t} = T_{\Sigma_{t'}}$, then for each solution Φ , Φ induces the same initial data on Σ_t and $\Sigma_{t'}$ — i.e., each solution is periodic with period $|t - t'|$. So if $T_{\Sigma_t} = T_{\Sigma_{t'}}$ for each Σ_t and $\Sigma_{t'}$, then every solution would have to be a constant function on V .

The maps $T_{\Sigma_t} : \mathcal{S} \rightarrow \mathcal{I}$ establish a simple relationship between our Hamiltonians $H : \mathcal{S} \rightarrow \mathbb{R}$ and $h : \mathcal{I} \rightarrow \mathbb{R} : h = H \circ T_{\Sigma_t}^{-1}$ (we could have taken this as our definition of h).

Together Ω and H determine the flow on \mathcal{S} that implements time translation at the level of solutions while together ω and h determine the flow on \mathcal{I} that implements time evolution of initial data. Since any T_{Σ} relates Ω and ω on the one hand, and H and h on the other, one would hope that it would also intertwine the group actions corresponding to these flows. This is indeed the case. Let us write $t \cdot_{\mathcal{S}} \Phi$ for the solution that results when we time-translate the solution Φ by t units and let us write $t \cdot_{\mathcal{I}} (\phi, \pi)$ for the state that initial data set (ϕ, π) evolves into after t units of time. Then we find that $t \cdot_{\mathcal{I}} T_{\Sigma}(\Phi) = T_{\Sigma}(t \cdot_{\mathcal{S}} \Phi)$.⁸³

Relative to a slicing, each solution Φ on V corresponds to a curve $(\phi(t), \pi(t))$ in the space of initial data, with $(\phi(t), \pi(t)) := T_{\Sigma_t}(\Phi)$. And a curve of this form

⁸²Here we use the fact that π takes values in W^* while $\dot{\phi}$ takes values in W ; and we rely on the natural measure induced by the geometry of V to allow us to treat $L(\Phi)$ as a function rather than an n -form.

⁸³That is, each T_{Σ} is equivariant for the \mathbb{R} -actions $\cdot_{\mathcal{S}}$ and $\cdot_{\mathcal{I}}$.

is always a dynamical trajectory in \mathcal{I} (i.e., an integral curve of the flow generating time evolution on \mathcal{I}). Conversely, a dynamical trajectory $(\phi(t), \pi(t))$ in \mathcal{I} determines a unique solution $\Phi := T_{\Sigma_0}^{-1}(\phi(0), \pi(0))$ — and this solution can be viewed as the result of laying down the instantaneous field configurations $\phi(t)$ on the instants Σ_t in the slicing.

5.4 Time and Change

Change consists in a single object having a given property at a given time and a distinct and incompatible property at a different time. Within the Lagrangian approach, it is easy enough to draw a distinction between those solutions that represent change and those that do not: the changeless solutions are those which are invariant under the action of a group of time translations. Correspondingly, we will say that a dynamical trajectory in the space of initial data represents a changeless reality when the corresponding solution on V is invariant under some time translation group.⁸⁴

This much is entirely straightforward. But it is worth pausing and thinking about how change is represented at the level at which physical quantities are represented by functions on \mathcal{S} and \mathcal{I} . In the case of quantities defined on the space of initial data, the story is straightforward. Let $f \in C^\infty(\mathcal{I})$ correspond to some determinable physical property of instantaneous states. Then if (ϕ_0, π_0) evolves into (ϕ_1, π_1) and $f(\phi_0, \pi_0) \neq f(\phi_1, \pi_1)$ then the solution including these states manifests change with respect to the property represented by f .⁸⁵ And we can of course go on to ask, e.g., about the rate of change of f along a dynamical trajectory.

But how should we phrase this in terms of functions defined on the space of solutions?

Suppose that we are interested in the quantity that measures the volume of the spatial region on which a given field takes on non-zero values. While such a quantity is represented within the Hamiltonian framework by a function $f : \mathcal{I} \rightarrow \mathbb{R}$, there is no function on the space of solutions that can be identified with this quantity — for such functions assign values to entire physically possible histories, and thus cannot represent quantities that take on different values at different

⁸⁴Naively, one might think that a dynamical trajectory in the space of initial data should count as representing a changeless reality only if it is constant — that is, if the system is represented as being in the same instantaneous state at each instant of time. But this would be a mistake. Consider a well-behaved theory set in Minkowski spacetime, and let Φ be solution invariant under the notion of time translation associated with inertial frame A but not invariant under that corresponding to inertial frame B . Surely this counts as changeless — and ought to whether we pass to the Hamiltonian picture via a slicing adapted to frame A (which leads to a dynamical trajectory according to which the state of the system is constant) or via a slicing adapted to frame B (which leads to a picture in which the state undergoes nontrivial evolution).

⁸⁵Even if Φ represents a state of affairs in Minkowski spacetime, changeless in virtue of being invariant under the notion of time translation associated with inertial frame A , it may still represent some physical quantities as undergoing change — such as the location of the centre of mass of a system relative to inertial frame B .

instants within a history (or rather, they cannot do so in the same direct way that functions on the space of initial data can).

However: intuitively, for each instant $\Sigma \subset V$ there is a function $f_\Sigma : \mathcal{S} \rightarrow \mathbb{R}$ such that $f_\Sigma(\Phi)$ is the volume of the support of our field on Σ in the solution Φ . So it is tempting to say that our chosen quantity is represented as exhibiting change in a solution Φ if $f_\Sigma(\Phi) \neq f_{\Sigma'}(\Phi)$ for instants $\Sigma, \Sigma' \subset V$, and that in order to speak of the rate of change of our quantity we need to consider a parameterized family Σ_t of instants, and calculate $\frac{d}{dt} f_{\Sigma_t}(\Phi)$.⁸⁶

Of course, in the present context, it makes sense to employ our preferred slicing in setting up this framework.⁸⁷ For each instant Σ_t in our slicing we have a symplectic isomorphism $T_{\Sigma_t} : \mathcal{S} \rightarrow \mathcal{I}$. If $f : \mathcal{I} \rightarrow \mathbb{R}$ is the function on the space of initial data that represents the quantity of interest, then $f_t := f \circ T_{\Sigma_t}$ is the desired function on the space of solutions that assigns to a solution Φ the value of f on the initial data that Φ induces on Σ_t . So each slicing σ determines a one-parameter family of functions on \mathcal{S} that encodes the instantaneous values of our chosen physical quantity relative to the instants in σ . So relative to a choice of slicing, it makes sense to ask whether this quantity undergoes change, what the rate of change is, and so on.

REMARK 31 (An Alternative Approach to Constructing $\{f_t\}$). In the present setting, rather than relying on our entire one-parameter family of isomorphisms, $\{T_{\Sigma_t}\}$, to set up our one-parameter family of functions $\{f_t\}$, we could have used Σ_0 to construct f_0 then used our dynamical time translation group to define

$$f_{-t}(\Phi) := f_0(t \cdot_{\mathcal{S}} \Phi).$$

5.5 Overview

We have seen that if we put in place a number of very strong assumptions, we get in return a very clear picture of the representation of time and change. The assumptions are: that our equations of motion, Δ , are second-order; that these equations have ideal existence and uniqueness properties and they derive from a Lagrangian, L , that has a dynamical symmetry group, τ , that arises from time translation group, $\bar{\tau}$, on our spacetime, V ; and that we have chosen a slicing σ of V that is adapted to $\bar{\tau}$.

Lagrangian Picture. The space of solutions, (\mathcal{S}, Ω) , is a symplectic space. The function, $H : \mathcal{S} \rightarrow \mathbb{R}$, that assigns to each solution the total instantaneous

⁸⁶For this suggestion, see, e.g., [Rovelli, 1991].

⁸⁷Otherwise we can run into trouble. Consider a Φ defined on Minkowski spacetime such that for each inertial observer the spatial volume of the region in which the field is nonzero is constant in time. Because of length contraction, relatively moving inertial observers will assign different values to this volume. So if we choose Σ and Σ' belonging to slicings corresponding to distinct inertial frames, then we find that $f_\Sigma(\Phi) \neq f_{\Sigma'}(\Phi)$ even though Φ is changeless according to each inertial observer.

energy relative to $\bar{\tau}$ is the symplectic generator of τ (and also the Noether conserved quantity associated with it).

Hamiltonian Picture. We are able to construct a Hamiltonian version of our theory: a symplectic space of initial data (\mathcal{I}, ω) equipped with a Hamiltonian $h : \mathcal{I} \rightarrow \mathbb{R}$ that generates the dynamics of the theory. The dynamics is encoded in an \mathbb{R} -action on \mathcal{I} that implements time evolution.

Relation between the Pictures. To each instant Σ in our slicing corresponds the symplectic isomorphism $T_\Sigma : \mathcal{S} \rightarrow \mathcal{I}$, that maps a solution Φ to the initial data that it induces on Σ . Each such T_Σ relates H and h on the one hand and Ω and ω on the other — and intertwines the action of the group implementing time translation on \mathcal{S} with the action of the group implementing time evolution on \mathcal{I} .

Time. Time, in one of its facets, is represented in this scheme by three \mathbb{R} -actions: the action via symmetries on V that implements time translation, the symplectic action implementing time translation on \mathcal{S} , and the symplectic action implementing time evolution on \mathcal{I} . Note: in some spacetimes there will only be one notion of time translation, in others there will be many.

Change. In the Lagrangian picture, changelessness is represented in a straightforward way — some solutions are invariant under a time translation group of their underlying spacetime. So change can be characterized as the absence of changelessness and the definition can then be translated into the language of the Hamiltonian approach. When it comes to representing change of given physical quantities via the behaviour of functions on the space of initial data and the space of solutions, things become a bit more interesting. Here it is the Hamiltonian picture that underwrites a straightforward approach: one finds the function on the space of initial data corresponding to the quantity of interest, and examines its behaviour as the state evolves. On the Lagrangian side, things are more complicated. No function on the space of solutions can directly represent a changeable quantity. But by employing the slicing-dependent correspondence between the two pictures one can find a one-parameter family of functions on the space of solutions, each of which describes the value of the quantity along a distinct instant from the slicing. One can use this one-parameter family to define the rate of change of the quantity; and so on.

6 COMPLICATIONS

The account of the previous section was underwritten by several very strong assumptions. I now want to consider the effect on the picture developed above if one or another of these assumptions is dropped. My strategy is to leave untouched the assumptions that we need to in order construct a Hamiltonian picture of the sort

developed above — that the equations of motion be second-order and that spacetime have enough geometrical structure to support slicings — and to consider the effect of dropping the assumptions: (i) that solutions are defined globally in time; (ii) that there is a unique maximal solution consistent with any initial data set; (iii) that the Lagrangian admits a dynamical time translation group that arises from a time translation group on spacetime. I will in this section consider only the effect of dropping one of (i)–(iii) at a time — in the next section I will turn to general relativity, which is a theory in which (i)–(iii) fail, as does the assumption that spacetime has enough solution-independent geometry to support slicings.

In briefest sketch, we find that:

1. If we drop the assumption that solutions exist globally in time, then time evolution is no longer implemented by an \mathbb{R} -action on \mathcal{I} , and \mathcal{S} and \mathcal{I} are no longer symplectically isomorphic. But time evolution is implemented by a sort of local and infinitesimal counterpart of an \mathbb{R} -action and \mathcal{S} and \mathcal{I} are *locally* symplectically isomorphic. Overall, only small changes are required in the picture of the representation of time and change developed above.
2. If we drop the assumption that specifying initial data suffices to determine a unique solution, even locally in time, by considering the (broad and important) class of theories whose Lagrangian and Hamiltonian versions exhibit gauge freedom, then the space of solutions and the space of initial data are presymplectic spaces that are not isomorphic (even locally). Furthermore, time evolution will no longer be implemented by a one-parameter group, but by a gauge equivalence class of such groups. Difficulties also appear on the Lagrangian side. The problem appears to be that theories of this type feature nonphysical variables. The remedy is reduction — the reduced space of solutions and the reduced space of initial data are symplectic and isomorphic. Much of the picture of the representation of time and change reappears at the reduced level.
3. If we drop the assumption that our Lagrangian is time-translation invariant, then we have to make do with time-dependent Lagrangian and Hamiltonian theories. Here the space of solutions and the space of initial data will be symplectic spaces, and will be isomorphic. But we no longer have time translation of solutions as a symmetry on the Lagrangian side, nor time evolution as a symmetry on the Hamiltonian side. Still, we are able to construct in the usual way a slicing-dependent one-parameter family of isomorphisms between the space of solutions and the space of initial data, and this allows us to reconstruct much of the familiar picture of the representation of time and change.

6.1 *Singular Dynamics*

Let us suppose that the condition of global existence of solutions fails for our equations of motion — there exist initial data sets that cannot be extended to

solutions defined on all of V . But let us continue to suppose that our theory is otherwise well-behaved: our spacetime V has enough structure to support slicings; our equations Δ are second-order and have unique solutions; and our Lagrangian, L , admits a dynamical time translation group, τ , induced by a time translation group $\bar{\tau}$ on V . Then, at least for the sorts of cases that arise in physics, we can expect to find the following.

Lagrangian Picture. The space of solutions, (\mathcal{S}, ω) , is a symplectic manifold. The dynamical time translation group, τ , acts on \mathcal{S} in the usual way: each element of the group time-translates each solution by some given amount.⁸⁸ τ is generated, via Ω , by the Hamiltonian function, $H : \mathcal{S} \rightarrow \mathbb{R}$ that assigns to a solution the instantaneous energy of that solution.

Hamiltonian Picture. We can construct a Hamiltonian picture as above: given a time translation group $\bar{\tau}$, an adapted slicing σ , a solution Φ , and an instant Σ , we can construct the initial data that Φ induces on Σ relative to σ , and use σ to pull this back to our abstract instant S . We can then construct the space of initial data, \mathcal{I} , with its canonical symplectic form ω , use our Lagrangian to define a Hamiltonian, $h : \mathcal{I} \rightarrow \mathbb{R}$, and study the resulting dynamics. The essential novelty is that because some solutions have limited temporal domains of definition, one finds that the vector field on (\mathcal{I}, ω) generated by the h is incomplete — it has integral curves that are defined only on a subset of \mathbb{R} . So time evolution is not represented by an \mathbb{R} -action on \mathcal{I} : in general it does not make sense to ask of a given point in the space of initial data what state it will evolve to at arbitrarily late times. However, the vector field generated by the Hamiltonian, which as usual encodes the dynamics, can be thought of as a sort infinitesimal generator of a locally defined action of \mathbb{R} on the space of initial data — in particular, if it makes sense to speak of data set x evolving into data set y after t units of time, then we find that the map that sends a state to the state t units of time later is a symplectic (and Hamiltonian-preserving) map between sufficiently small neighbourhoods of x and sufficiently small neighbourhoods of y .

Relation between the Pictures. As above, for each instant, Σ , in our slicing we can define $T_\Sigma(\Phi)$ to be the pullback to the abstract instant S of the initial data that the solution Φ induces on Σ . But now each T_Σ is only partially defined as a function from \mathcal{S} to \mathcal{I} (since the value of $T_\Sigma(\Phi)$ is undefined when Φ is not defined on Σ). Nonetheless, each such T_Σ is a symplectic isomorphism between its domain of definition in \mathcal{S} and \mathcal{I} .⁸⁹ As usual, we get a distinct such map for each instant we choose.

⁸⁸Of course, if a solution is not defined for all time, then its domain of definition will differ from that of its time-translate in the obvious way.

⁸⁹So, intuitively, the space of solutions is bigger than the space of initial data — we can find natural isomorphisms between the space of initial data and subspaces of the space of solutions.

Time. The representation of time becomes a bit more complicated in the present context: to each notion of time translation on spacetime corresponds a nice symmetry on the space of solutions — and a merely infinitesimal symmetry on the space of initial data.

Change. We can still represent changeable properties by functions on \mathcal{I} , and determine whether a given dynamical trajectory represents a change of such properties by studying the behaviour of the corresponding function along the trajectory. Despite the failure of global isomorphism between the space of solutions and the space of initial data, we find that a choice of slicing yields a one-parameter family of local isomorphisms, $\{T_{\Sigma_t}\}$, between \mathcal{I} and subspaces of \mathcal{S} . Given a function f on the space of initial data corresponding to a quantity of interest, the family $\{T_{\Sigma_t}\}$ can be used to construct a one-parameter family of partially-defined functions $\{f_t\}$ on \mathcal{S} that correspond to the given changeable physical quantity. So the representation of change in this case is much the same as in the case in which we have global existence of solutions.

The real novelty here is the lack of a global isomorphism between the space of solutions and the space of initial data. The phenomenon can be well-illustrated by simple classical mechanical examples.

EXAMPLE 32 (The Kepler Problem). Consider a point-particle of mass m moving in the x - y plane subject to the gravitational influence of a point-particle of unit mass fixed at the origin.⁹⁰ Here our spacetime V will be \mathbb{R} and the space W of field values will be the space $Q = \{(x, y)\}$ of possible positions of the moving particle. The Lagrangian is $L = \frac{1}{2}(\dot{x}^2 + \dot{y}^2) + \frac{m}{r}$, where $r^2 := x^2 + y^2$; the corresponding Hamiltonian is $H = \frac{1}{2}(\dot{x}^2 + \dot{y}^2) - \frac{1}{r}$. In order for L and H to be well-defined, we have to limit the location of the body to points in $Q := \mathbb{R}^2 \setminus \{(0, 0)\}$. We restrict attention to the case where $H < 0$. This is the case of orbits bounded in space — so, in particular, we rule out parabolic and hyperbolic motions.

We find that there are two types of solutions. (i) Regular solutions, in which the particle has non-zero angular momentum, are periodic and defined for all values of t ; they represent the particle as moving along an ellipse that has the origin as a focus. (ii) Singular solutions, in which the particle has vanishing angular momentum, are defined only for $t_0 < t < t_0 + 2\varepsilon$; they represent the particle as being ejected from the origin at time t_0 (i.e., $|r(t)| \rightarrow 0$ as $t \rightarrow t_0$ from above), travelling outward along a ray from the origin with decreasing speed until reaching to a stop at time $t_0 + \varepsilon$, then falling back into the origin along the same ray, with $|r(t)| \rightarrow 0$ as $t \rightarrow t_0 + 2\varepsilon$ from below.

The space of solutions is topologically pathological. Let $\Phi(t)$ be a singular solution defined for $t \in (t_0, t_0 + 2\varepsilon)$. Let $\Lambda \subset Q$ be the line segment along which the particle moves according to Φ . It is possible to construct a sequence $\{\Phi_k\}$ of

⁹⁰For the structure of the space of solutions of the Kepler problem, see [Woodhouse, 1991, §2.3] and [Marco, 1990b].

regular solutions with the following features: each Φ_k has the same energy as Φ — it follows that each Φ_k represents the particle as moving periodically along an ellipse E_k with period 2ε ; for each k , E_k is oriented so that the segment joining its foci is included in Λ ; as $k \rightarrow \infty$, the eccentricity of E_k goes to infinity — so that $E_k \rightarrow \Lambda$ as $k \rightarrow \infty$. It follows that Φ is a limit of the sequence $\{\Phi_k\}$.⁹¹ But now consider the solution $\Phi'(t) := \Phi(t + 2\varepsilon)$. This is a singular solution, defined for $t \in (t_0 - 2\varepsilon, t_0)$, that represents the particle as being emitted at time $t_0 - 2\varepsilon$, moving along Λ , and then being absorbed at time t_0 . Φ' is equally a limit of $\{\Phi_k\}$. Indeed, we can generate an infinite number of limits of $\{\Phi_k\}$ by temporally translating Φ by multiples of 2ε .

Since we can find a sequence in \mathcal{S} that converges to more than one limit point, \mathcal{S} is not Hausdorff.⁹² But, of course, the space of initial data for the Kepler problem is just T^*Q — which is Hausdorff. So the spaces are certainly not isomorphic!

Non-Hausdorff manifolds can be quite wild. But there are also relatively tame examples, such as the following. Let X be the result of excising the origin from the real line and adding in its place two new objects, a and b ; a subset of X is an open ball if it coincides with an open interval in \mathbb{R} that does not contain the origin, or if it arises when one takes an open interval of real numbers containing 0 and replaces 0 by one of $\{a, b\}$. We endow X with a topology by declaring that any union of open balls in X is an open set. X is a manifold according to our present standards. But it is not Hausdorff, since every neighbourhood of a overlaps with a neighbourhood of b — and, of course, a sequence like $\{\frac{1}{k}\}$ converges to both a and b .

More generally, we can construct a non-Hausdorff manifold $X_j^{n,m}$ by taking m copies of \mathbb{R}^n and identifying them everywhere except on a given j -dimensional hyperplane through the origin ($1 < m \leq \infty$, $1 \leq n < \infty$, and $0 \leq j < n$).⁹³ The space of solutions corresponding to any fixed negative value of energy in the planar Kepler problem is assembled out of copies of $X_1^{3,\infty}$.

EXAMPLE 33 (Singularities of the n -Body Problem). For n particles in \mathbb{R}^3 the space of possible particle configurations is \mathbb{R}^{3n} . But this space includes collisions — and the potential energy for the n -body problem is singular at such points. So, as before, we let $\Delta := \{q \in \mathbb{R}^{3n} : \mathbf{q}_i = \mathbf{q}_j \text{ for some } i \neq j\}$ and let Q be \mathbb{R}^{3n}/Δ then take as our space of initial data $T^*Q = \{(q, p) \in T^*\mathbb{R}^{3n} : q \notin \Delta\}$.

We pose the initial data (q, p) at time $t = 0$. We know that this determines a maximal dynamical trajectory $t \mapsto (q(t), p(t))$, defined on an interval $[0, t^*)$, with $0 < t^* \leq \infty$ (the corresponding story for negative times is, of course, the same).

⁹¹The topology on the space of solutions can be constructed as follows. For each $t \in \mathbb{R}$, take the position and velocity of the particle at a given time as coordinates on the space of solutions, and construct the set of open balls relative to these coordinates. Now take the union of these sets as t varies. The result is a basis for the topology we seek.

⁹²Recall that a topological space X is *Hausdorff* if for any $x, y \in X$ it is possible to find disjoint open $U, V \subset X$ with $x \in U$ and $y \in V$. A sequence in a Hausdorff space has at most one limit.

⁹³The example of the preceding paragraph is $X_0^{1,2}$.

Clearly it is possible to choose (q, p) so that t^* is finite — if we let $p = 0$ for $n > 1$, for instance, the system is going to collapse and a collision will occur. Let us call our dynamical trajectory *singular* if $t^* < \infty$. It can be proved that if $t^* < \infty$, then as $t \rightarrow t^*$, $q(t) \rightarrow \Delta$, in the sense that $\lim_{t \rightarrow t^*} \min_{1 \leq i < j \leq n} r_{ij} = 0$. Let us say that a singular trajectory ends in a *collision* if there is a point $q_0 \in \Delta$ such that $\lim_{t \rightarrow t^*} q(t) = q_0$; otherwise, we say that it ends in a *pseudocollision*.

Consider the following cases.⁹⁴

n = 1. This is the case of a single free particle. The dynamics is non-singular.

n = 2. This is the Kepler problem.⁹⁵ The only singularities are collision singularities. These occur if and only if the angular momentum of the system vanishes.

Famously, these singularities can be regularized.⁹⁶

This is clear enough physically: one simply imposes the condition that any collisions that occur are elastic. This allows one to sew together a solution which ends with a collision at time t_0 with one, that, intuitively begins at time t_0 with the particles having interchanged their velocities. Continuing in this way, one constructs continuous and piecewise analytic solutions of eternal temporal extent. Because collision solutions are now infinite in temporal extent, the space of solutions, in this new sense, is isomorphic to the extended space of initial data that includes the collision states that lie in $T^*\Delta$ (let us interpret such states as representing the velocity that the particles will have when next emitted).

Mathematically, there are a number of underpinnings that can be given to this procedure.⁹⁷ An older one proceeds in terms of *analytic continuation* — thinking of the original collision solution as a complex function, one asks whether there is any analytic continuation of this function past the time of collision. Under a more modern approach, one looks for a way of continuing singular solutions that preserves the continuous dependence of evolution on initial data.⁹⁸ In the case of the two-body problem, either approach vindicates

⁹⁴For surveys of the singularities of the n -body problem, see [Diacu, 1992] and [Diacu, 2002]; for a popular treatment, [Diacu and Holmes, 1996, Chapter 3].

⁹⁵Begin with the two-body problem. Restrict attention to the plane of motion of the particles; choose a frame in which the centre of mass of the two-body system is at rest at the origin and denote the positions of the bodies as \vec{q}_1 and \vec{q}_2 . Obviously if we know $\vec{r} = \vec{q}_2 - \vec{q}_1$ then we know the positions of both particles (since we know their masses and the location of their centre of mass). Now note that the equation of motion for \vec{r} is that for a single particle moving in an gravitational potential around the origin, if we take the origin to have unit mass, and the moving particle to have mass $m = \frac{m_1 m_2}{m_1 + m_2}$.

⁹⁶For various approaches to the regularization of two-body collisions, see [Souriau, 1982], [Marco, 1990a] and [Cushman and Bates, 1997, §II.3].

⁹⁷See [McGehee, 1975] for these.

⁹⁸More precisely, one excises from $T^*\mathbb{R}^{6n}$ an open set with compact closure that contains a collision; the boundary of this set falls into two pieces, corresponding to initial data of trajectories entering the set and initial data of trajectories leaving the set; evolution gives a diffeomorphism

cates extension of singular solutions by elastic collisions as the unique tenable method of extension.

n = 3. Singular trajectories end in collisions. Collisions involving only two bodies can be regularized as elastic collisions. But some three body collisions are non-regularizable (according to any of several criteria).⁹⁹ Such three-body collisions are complex, and presumably make it difficult to determine the topology of the space of solutions — so in this case, unlike the $n = 2$ case, one does not have a clear picture of the relation between the global structure of the space of solutions and the structure of the space of initial data.

n ≥ 4. For $n ≥ 4$, as usual, singular trajectories can end in collisions: two-body collisions are regularizable; but at least some collisions involving larger numbers of particles are not. Furthermore, for $n > 4$ it is known that pseudocollisions can also occur — so it would appear to be more difficult then ever to determine the topology of the space of solutions.¹⁰⁰

REMARK 34 (Quantization of Singular Systems). When the space of solutions and the space of initial data are isomorphic, it is, of course, a matter of indifference which space one takes as the starting point for quantization. When dynamics is singular and these spaces are no longer isomorphic one faces a real choice. And the choice is not entirely pleasant — one has to choose between the space of initial data, on which the dynamics is implemented by an incomplete vector field, and the space of solutions, which one expects to have a complicated and pathological topology. Presumably there is no guarantee that the two approaches always lead to the same quantizations in the domain of singular dynamics.¹⁰¹

6.2 Gauge Freedom

We next want to consider what happens when we drop the assumption that specifying initial data suffices to determine a unique maximal solution to our equations of motion. To this end, we are going to assume that our equations of motion *under-determine* the behaviour of the field, in the radical sense that for given

from the subset of the former corresponding to non-singular solutions to the subset of the latter corresponding to non-singular solutions; one asks whether this can be extended to a diffeomorphism of one whole piece to the other.

⁹⁹See [McGehee, 1975].

¹⁰⁰ See [Saari and Xia, 1995]. The question is open for $n = 4$; but see [Gerver, 2003] for a possible example. Pseudocollisions require that the positions of at least some of the particles become unbounded as $t \rightarrow t^*$ — by exploiting arbitrarily great conversions of potential energy into kinetic energy, these particles escape to infinity in a finite time. As emphasized by Earman ([1986, Chapter III] and [this volume]), this means that pseudocollisions involve a rather radical and surprising failure of determinism — which is most dramatic when one considers the time reverse of such a process, in which particles not originally present anywhere in space suddenly swoop in from infinity.

¹⁰¹For approaches to the quantization of systems with singular dynamics see, e.g., [Gotay and Demaret, 1983] and [Landsman, 1998].

initial data the general solution consistent with that data contains at least one arbitrary function of the full set of independent variables.¹⁰² There is a wide class of physical theories whose equations exhibit this *prima facie* pathological behaviour — including, most importantly, Maxwell’s theory of electromagnetism, general relativity, and their generalizations.

In this subsection I will first sketch a little bit of the theory of Lagrangian treatments of such theories without making any special assumptions about time translation invariance, the global existence of solutions, or the structure of spacetime. These further assumptions will later be brought into play, and will underwrite a consideration of the Hamiltonian form of a theory that is well-behaved except in possessing under-determined dynamics, and of the representation of time and change in such theories. This discussion will be followed by three examples.

Let us begin by introducing the notion of a family of gauge symmetries of a Lagrangian theory. Recall that a group, G , acting on the space, \mathcal{K} , of kinematically possible fields is a group of variational symmetries of a Lagrangian, L , defined on \mathcal{K} if the action of G is appropriately local and leaves the variational problem of L invariant.¹⁰³ We call a group, G , of variational symmetries of (\mathcal{K}, L) a group of *gauge symmetries* if it can be parameterized in a natural way by a family of arbitrary functions on spacetime.¹⁰⁴ Roughly speaking, each function on spacetime generates a Noether group of symmetries of (\mathcal{K}, L) — a one-parameter group of (suitably local) symmetries of the variational problem of L .¹⁰⁵ Since the set of functions on spacetime is infinite-dimensional, any group of local symmetries of a theory is infinite-dimensional.

Most familiar groups of symmetries of physical theories — the group of isometries of a spacetime with non-trivial geometry, the group that acts by changing the phase of the one-particle wavefunction by the same factor at each spacetime point, etc. — are finite-dimensional, and hence not do not count as groups of gauge symmetries in the present sense.

It is easy to see that the equations of motion of a Lagrangian theory admitting such a group of gauge symmetries under-determine solutions to the theory. Let

¹⁰²On the relevant notion of under-determined equations of motion, see, e.g., [Olver, 1993, 170–172, 175, 342–346, and 377].

¹⁰³For a more precise definition, see [Zuckerman, 1987, 274].

¹⁰⁴Let us be more precise. First, let Y be a vector space, and let Γ be a space of functions from V to Y (more generally, let Γ be a space of sections of some vector bundle $E \rightarrow V$). We assume that Γ includes all smooth, compactly supported maps from V to W , but leave open the precise boundary conditions, smoothness conditions, etc., required to characterize Γ . (Special care regarding boundary conditions is required when Γ contains elements with noncompact support).

Now we define a group of *gauge symmetries* parameterized by Γ as a pair of linear and local maps, $\varepsilon \mapsto X_\varepsilon$ and $\varepsilon \mapsto R_\varepsilon$ sending elements ε of Γ to local vector fields on S and to elements of $\Omega_{loc}^{n-1,0}(V \times \mathcal{K})$, respectively, such that $\partial L(\Phi, X_\varepsilon(\Phi)) = DR_\varepsilon(\Phi)$ for all $\Phi \in S$ and $\varepsilon \in \Gamma$. So each $\varepsilon \in \Gamma$ is associated with an infinitesimal generator of a Noether group for L (cf. fn. 71 above).

¹⁰⁵A bit more precisely: the discussion of the previous footnote shows that the map $\varepsilon \mapsto (X_\varepsilon, R_\varepsilon)$ is a map from Γ to the set of generators of Noether groups of (\mathcal{K}, L) ; in fact, the image of this map will be infinite-dimensional in nontrivial examples, although it may have a nontrivial kernel (in example 38 below, constant functions all generate the same (trivial) Noether group).

ε be a function on spacetime that vanishes everywhere but on some compact set $U \subset V$; if we allow the corresponding Noether group $\xi = \{g_t\}$ to act on a solution Φ , then for $t \neq 0$ the resulting solutions $\Phi_t = g_t \cdot \Phi$ will agree with Φ outside of U , but in general disagree with Φ inside U . Thus if we choose an instant $\Sigma \subset V$ which does not intersect U , we find that Φ and Φ_t induce the same initial data on Σ , but differ globally — so uniqueness fails for the equations of motion of the theory.

Recall from section 3.3 that a presymplectic form is a degenerate closed two-form, and that the imposition of such a form on a space serves to partition the space by submanifolds called gauge orbits. An argument very similar to that of the previous paragraph shows that if L admits a group of gauge symmetries, then the form Ω that L induces on the space of solutions is presymplectic, and that the corresponding gauge orbits are such that two solutions belong to the same gauge orbit if and only if they are related by an element of the group of gauge symmetries of L .¹⁰⁶ So gauge symmetries of L are gauge transformations of (\mathcal{S}, Ω) , in the sense stipulated in section 3.3 above — they preserve the gauge orbits of the space of solutions.

It follows from general facts about presymplectic forms that if a function on \mathcal{S} generates a one-parameter group of gauge symmetries, then that function is a constant function. In particular: the Noether conserved quantity, $Q_\xi : \mathcal{S} \rightarrow \mathbb{R}$, associated with a one-parameter group of gauge symmetries, ξ , must be a constant function.¹⁰⁷ Such conserved quantities are trivial, in the sense that they do not provide any means to distinguish between physically distinct solutions.

So much we can say about any Lagrangian theory admitting a group of gauge symmetries. Let us now specialize to the case where our equations of motion, Δ , are second-order, our spacetime, V , has enough geometrical structure to admit slicings, solutions exist globally in time, and our Lagrangian, L , admits a dynamical time translation group, τ . With these further assumptions in place, we can investigate the implications that giving up on local uniqueness of solutions has for the picture of time and change developed in section 5 above. We find the following.

Lagrangian Picture. We are assuming that we have a notion of time translation arising out of the structure of our background spacetime V . This notion gives rise, in the usual way, to a dynamical time translation group, τ . The corresponding conserved quantity is the usual Hamiltonian, H , which assigns to each solution its instantaneous energy. So far so good. But now recall from the discussion of section 3.3 above that in the setting of a presymplectic space, if a function generates a given one-parameter family of transformations of the space, then it also generates all one-parameter families of trans-

¹⁰⁶Recall that in fn. 61 above, Ω was defined as the integral of a certain object over an arbitrary instant $\Sigma \subset V$. If we consider an infinitesimal local symmetry X_ε which has no effect on solutions along Σ , then Ω will not see X_ε — i.e., $X_\varepsilon(\Phi)$ will be a null vector at each $\Phi \in \mathcal{S}$. See [Deligne and Freed, 1999, §2.5] and [Woodhouse, 1991, 145].

¹⁰⁷In fact, it will be the zero function, because the Noether current J_ξ will be exact as an $(n-1)$ -form on V . See [Zuckerman, 1987, 274].

formation gauge equivalent to the given one. In the present case, this means that in addition to the dynamical time translation group, τ , H generates all one-parameter groups of transformation of (\mathcal{S}, Ω) that agree up to gauge with τ .

Hamiltonian Picture. Having fixed a notion of time translation in spacetime and an associated slicing of spacetime into instants, we can proceed as usual to construct the space of initial data that arise when the configuration and momentum variables of the field are restricted to an arbitrary instant in our slicing.¹⁰⁸ As in the well-behaved case, given an instant $\Sigma \subset V$ of our slicing and a solution Φ , we can construct a corresponding initial data set (ϕ, π) on our abstract instant S , by pulling back to S the initial data that Φ induces on Σ . In the well-behaved case, we found that the space of initial data had the structure $T^*\mathcal{Q}$, where \mathcal{Q} was the space of all ϕ that arise as instantaneous field configurations by restricting solutions to instants. In the present case we find that the (ϕ, π) that arise as initial data sets form a subspace of $T^*\mathcal{Q}$ (where \mathcal{Q} is again the space of all ϕ that arise as restrictions of solutions to instants).¹⁰⁹ In addition, we may also find that in order to construct consistent dynamics, we need to further restrict admissible initial data. The upshot is that we take as our space of initial data a subspace $\mathcal{I} \subset T^*\mathcal{Q}$. \mathcal{I} comes equipped with a natural geometric structure: the cotangent bundle $T^*\mathcal{Q}$ comes equipped with its canonical symplectic form (see example 7 above); the restriction of this form to \mathcal{I} yields a presymplectic form ω . When all goes well, the gauge orbits determined by ω have the following structure: initial data sets (ϕ, π) and (ϕ', π') arising as the initial data induced on a given instant $\Sigma \subset V$ by solutions Φ and Φ' belong to the same gauge orbit in (\mathcal{I}, ω) if and only if Φ and Φ' belong to the same gauge orbit in (\mathcal{S}, Ω) .¹¹⁰ One can go on to define a Hamiltonian function, h , on (\mathcal{I}, ω) in the usual way. Of course, since (\mathcal{I}, ω) is a merely presymplectic space, h generates a whole gauge equivalence class of notions of dynamics (i.e., one-parameter groups of symmetries of (\mathcal{I}, ω)). Suppose that according to one such notion of dynamics, initial state x_0 evolves into state $x(t)$ at time t . Then although other notions of time evolution generated by h will in general disagree about what state x_0 evolves into at time t , they will all agree that the state at t lies in $[x(t)]$, the gauge orbit of $x(t)$.¹¹¹

¹⁰⁸On constructing the constrained Hamiltonian system corresponding to a given Lagrangian theory admitting gauge symmetries, see [Dirac, 2001], [Gotay *et al.*, 1978], and [Henneaux and Teitelboim, 1992]. For philosophical discussion, see [Earman, 2003] and [Wallace, 2003].

¹⁰⁹This is because so-called *first-class constraints* arise: it follows from the definition of the momenta, $p_i := \frac{\partial L}{\partial \dot{q}_i}$, that some components of the momenta are required to vanish identically.

¹¹⁰*Warning:* it is not difficult to construct (unphysical) examples in which this nice picture fails — see example 36 below.

¹¹¹In fact, for each point $y \in [x(t)]$, there is a notion of time evolution generated by h according to which x_0 evolves into y at time t .

Relation between the Pictures. As usual, for each instant Σ in our slicing we can define $T_\Sigma : \mathcal{S} \rightarrow \mathcal{I}$, the map that sends a solution to the (pullback to S of the) initial data that it induces on Σ . In the setting of section 5, these maps gave us isomorphisms between the space of solutions and the space of initial data. But in the presence of gauge symmetries, these maps are *not* isomorphisms — since the existence of gauge symmetries implies that many solutions induce the same initial data on any given Σ . The situation is most dramatic when we consider a theory with only finitely many degrees of freedom which admits gauge symmetries: for then the space of solutions will be infinite-dimensional while the space of initial data will be finite-dimensional (see example 37 below).¹¹² When all goes well, we get the following picture of the relation between solutions and dynamical trajectories in the space of initial data (holding fixed a notion of time translation and a slicing adapted to it).

1. Let Φ be a solution and let $x(t) = (\phi(t), \pi(t))$ be the curve in \mathcal{I} that arises by letting $x(t)$ be the initial data set that Φ induces on the instant $\Sigma_t \subset V$. Then $x(t)$ is a dynamical trajectory of the Hamiltonian version of the theory.
2. Given a dynamical trajectory, $x(t)$, of the Hamiltonian version of the theory we find that there is a unique solution $\Phi \in \mathcal{S}$ such that the curve in \mathcal{I} that corresponds to Φ in the sense of the preceding clause is just $x(t)$.
3. If $\Phi, \Phi' \in \mathcal{S}$ belong to the same gauge orbit in the space of solutions, then the corresponding dynamical trajectories, $x(t)$ and $x'(t)$ in \mathcal{I} agree up to gauge (in the sense that for each t , $x(t)$ and $x'(t)$ belong to the same gauge orbit in \mathcal{I}).
4. If dynamical trajectories $x(t)$ and $x'(t)$ in \mathcal{I} agree up to gauge, then the solutions $\Phi, \Phi' \in \mathcal{S}$ to which they correspond belong to the same gauge orbit in \mathcal{S} .

Time. Our notion of time translation lifts in a nice way from our spacetime V to the space of solutions, \mathcal{S} , where we get the usual representation of time via an \mathbb{R} -action. Even here there is an oddity: the Hamiltonian that generates this action also generates each \mathbb{R} -action gauge-equivalent to it. The situation is messier still in the space of initial data: given a notion of time translation on spacetime and a slicing adapted to that notion, we can construct a Hamiltonian picture; but in the presence of gauge symmetries, we find that there are many dynamical trajectories through each point in the space of initial data. In effect, our single notion of time translation in spacetime splits into a multitude of \mathbb{R} -actions on the space of initial data, each with equal claim to be implementing the dynamics of the theory.

¹¹²So in this case we see that no map from (\mathcal{S}, Ω) to (\mathcal{I}, ω) is an isomorphism; intuitively this is true for any theory admitting local symmetries.

Change. The evolution of arbitrary quantities under the dynamics defined on the space of initial data is indeterministic: if $x_0 = (\phi_0, \pi_0)$ is an initial data set, there will be distinct dynamical trajectories $x(t)$ and $x'(t)$ passing through x_0 at time $t = 0$; for an arbitrary function $f : \mathcal{I} \rightarrow \mathbf{R}$, we have no reason to expect that $f(x(t)) = f(x'(t))$ for $t \neq 0$; so fixing the state at time $t = 0$ does not suffice to determine the past and future values of the quantity represented by f . But since in this situation $x(t)$ and $x'(t)$ will agree about which gauge orbit the state lies in at each time, we find that the evolution of gauge-invariant quantities (those represented by functions on the space of initial data that are constant along gauge orbits) is fully deterministic — given the initial state, one can predict the value of such a quantity at all times. Furthermore, our slicing allows us to associate with each gauge-invariant function, f , on the space of initial data a one-parameter family $\{f_t\}$ of gauge-invariant functions on the space of solutions: let $f_t(\Phi)$ be the value that f takes on the initial data set that Φ induces on Σ_t . In this way we can represent change of gauge-invariant quantities via functions on the space of initial data or the space of solutions in the usual way.

This last point, especially, ought to arouse the suspicion that our theory, in the form currently under consideration, contains surplus structure. For while the theory has some quite disappointing features — ill-posed initial value problem, trivial conservation laws, a merely presymplectic geometric structure, failure of even local isomorphism between the space of solutions and the space of initial data — one finds that there is a large subset of physical quantities that behave just as the quantities of a well-behaved theory do. One naturally wonders whether there might be a well-behaved theory governing the behaviour of these quantities lurking somewhere in the background.

This sort of suspicion motivates the application to (\mathcal{S}, Ω) and (\mathcal{I}, ω) of the reduction procedure discussed in section 3.3 above. When all goes well, the following picture emerges: the reduced space of solutions (i.e., the space of gauge orbits of the space of solutions) and the reduced space of initial data (i.e., the space of gauge orbits of the space of initial data) are both symplectic spaces — and these reduced spaces are isomorphic.¹¹³ The Hamiltonian functions corresponding to time translation on the original space of solutions and time evolution on the original space of initial data project down to the reduced spaces. The resulting reduced Hamiltonians generates time translation and time evolution on their respective spaces.

In typical examples that arise in physics, one sees that the original theory's invariance under a group of gauge symmetries was in fact a sign that physically otiose variables had been included in the theory. Indeed: the fact that the original space of initial data is presymplectic with a symplectic reduced space indicates that within the original Hamiltonian formulation of the theory one can partition the set of variables parameterizing the original space of initial data into two classes, that we will call the class of *physically relevant variables* and the class of *physically*

¹¹³Example 36 below is an (unphysical) case where this isomorphism fails.

otiose variables; specifying the initial values of all variables suffices to determine the values for all times of the physically relevant variables while leaving wholly arbitrary the evolution of the physically otiose variables.¹¹⁴ At least locally the physically relevant variables can be taken to parameterize the reduced space of initial data. This provides a strong reason to think that the Hamiltonian theory defined on the reduced space of initial data gives a perspicuous representation of the physics under investigation, involving as it does exactly those quantities whose evolution is determined by the original theory. And this in turn provides good reason to think of the reduced space of solutions as representing possible histories of the system whose possible instantaneous states are represented by points in the space of initial data.¹¹⁵

REMARK 35 (Reduction and Determinism). Suppose that one is presented with a *prima facie* indeterministic theory, in which many future sequences of states are consistent with a given initial state. Then one could always construct a deterministic theory by simply identifying all of the futures consistent with a given state. As noted by Maudlin, it would be foolish to apply this strategy whenever one encountered an indeterministic theory: (i) general application of this strategy would render determinism true by fiat; and (ii) one would often end up embracing trivial or silly theories.¹¹⁶ For example: in Newtonian physics, the initial state in which space is empty of particles is consistent with a future in which space remains empty, and also with a future in which particles swoop in from infinity, then interact gravitationally for all future time (see fn. 100 above); to identify these futures — to view them as mere re-descriptions of a single physical possibility — would be absurd.

Now, reduction is a special case of the general strategy that Maudlin objects to. But since just about any wise course of action is a special case of a strategy that is in general foolish, this is not in itself an objection to reduction. We ought to check whether the complaints that Maudlin quite rightly registers against the general strategy redound to the discredit of the special case. I claim that they do not. (i) It is true that reducing theories with gauge symmetries converts *prima facie* indeterministic theories into deterministic ones. But this is unobjectionable: the sort of indeterminism that is a *prima facie* feature of a theory with gauge symmetries (namely, the existence of quantities whose evolution is wholly unconstrained by the initial state of the system) appears to be unphysical. (ii) For the

¹¹⁴In a presymplectic manifold satisfying suitable technical conditions, every point has a neighbourhood admitting a chart whose coordinates fall into two classes — those that parameterize gauge orbits and those that parameterize the directions transverse to the gauge orbits; see [Abraham and Marsden, 1978, Theorem 5.1.3]. In the space of initial data, it is natural to take the variables of the first type to be physically otiose and those of the second type to be physically relevant.

¹¹⁵Since this space arises by identifying solutions related by elements of the group of gauge symmetries of the theory, while the reduced space of initial data arises by identifying initial data that are induced on a given instant by solutions related by elements of the group of gauge symmetries of the theory.

¹¹⁶See [Maudlin, 2002, 6–8].

sort of theories that arise in physics, one does not have to fear that reduction will lead to a trivial or absurd result — in known cases, reduction carries one to a well-behaved symplectic space that is a suitable setting for a physical theory. Indeed, in such cases, it is (almost unanimously) agreed that the resulting symplectic space parameterizes the true degrees of freedom and provides the correct setting for the dynamics of the original theory.¹¹⁷

EXAMPLE 36 (A Pathological example). Before proceeding, it is important to emphasize that it is not hard to cook up simple (but unphysical) examples that do not follow the pattern sketched above for theories with gauge symmetries.¹¹⁸ Consider a particle moving in the x - y plane with Lagrangian $L = \frac{1}{2}e^y\dot{x}^2$. The corresponding Euler–Lagrange equations tell us that x is constant in time while the evolution of y is wholly arbitrary. So the space of solutions consists of pairs $(x_0, y(t))$ where $x_0 \in \mathbb{R}$ and $y(t) : \mathbb{R} \rightarrow \mathbb{R}$ an arbitrary smooth function; two solutions $(x_0, y(t))$ and $(x'_0, y'(t))$ belong to the same gauge orbit if and only if $x_0 = x'_0$. So the reduced space of solutions is just \mathbb{R} — which, having an odd number of dimensions, cannot carry a symplectic structure. On the Hamiltonian side one finds that the momentum conjugate to x and the momentum conjugate to y both have to vanish — which means that the space of initial data is $\mathbb{R}^2 = \{(x, y)\}$, with every point being gauge equivalent to every other.¹¹⁹ So the reduced space of initial data is a single point — which is not isomorphic to the reduced space of solutions.

EXAMPLE 37 (Particles on a Line). We consider two gravitating point-particles moving on a line. For simplicity, we choose units so that Newton’s constant is unity, assume that the particles have unit mass, and set aside worries about collisions and their regularization. We consider three theories of this system.

The Newtonian Theory. We denote the positions of the particles as q_1 and q_2 with $q_2 > q_1$. We interpret these as giving the positions of the particles relative to a frame at rest in absolute space. The Lagrangian for this system is $L = T - V$ where the kinetic energy is $T := \frac{1}{2}(\dot{q}_1^2 + \dot{q}_2^2)$ and the potential energy is $V = -\frac{1}{q_2 - q_1}$. The usual Newtonian equations of motion follow. It is helpful to consider a variant formulation of this theory. We define new configuration variables, $r_0 := \frac{1}{2}(q_1 + q_2)$ and $r_1 := q_2 - r_0 = \frac{1}{2}(q_2 - q_1)$ (so r_0 is the position of the centre of mass of the system, and r_1 is half the relative distance between the particles). In terms of these variables, our Lagrangian

¹¹⁷General relativity provides the sole instance in which there is any dissent from the consensus view; see [Kuchař, 1986] and [Kuchař, 1993]. This is also the case that Maudlin is concerned with — he, like Kuchař, worries that unreflective application of reduction to general relativity leads to absurd conclusions about time and has hampered conceptual progress in quantum gravity. Part of the burden of section 7 below is to show that no absurdities follow from the application of reduction in that case.

¹¹⁸For the following example, see [Henneaux and Teitelboim, 1992, §1.2.2]. For further discussion of such examples, see [Gotay, 1983].

¹¹⁹This is an example where one constraint arises directly from the definition of the momenta while the other is required in order to formulate consistent dynamics.

is $L(r_0, r_1, \dot{r}_0, \dot{r}_1) = \frac{1}{2}(\dot{r}_0^2 + \dot{r}_1^2) + \frac{1}{2r_1}$. The equations of motion tell us that r_0 is a linear function of time (since the centre of mass of an isolated system moves inertially) while $r_1(t)$ solves $\ddot{r}_1 = -\frac{1}{2r_1^2}$, and so describes the relative motion between the particles as they interact gravitationally.

The Leibnizean Theory. In this theory, space and motion are relative, and so the relative distance, r_1 , between the particles is the only configuration variable (or rather, r_1 is half the relative distance). The Lagrangian for the Leibnizean theory is $L'(r_1, \dot{r}_1) := \frac{1}{2}\dot{r}_1^2 + \frac{1}{2r_1}$. The equation of motion is $\ddot{r}_1 = -\frac{1}{2r_1^2}$. So the Leibnizean theory gives the same dynamics for the relative distances between the particles as the Newtonian theory.

The Semi-Leibnizean Theory. We take both r_0 and r_1 as configuration variables, and take as our Lagrangian $L''(r_0, r_1, \dot{r}_0, \dot{r}_1) := \frac{1}{2}\dot{r}_1^2 + \frac{1}{2r_1}$ (so L'' is a function of r_0 , r_1 , \dot{r}_0 , and \dot{r}_1 which happens to depend only on r_1 and \dot{r}_1). We apply the variational algorithm: as always, it leads to the conclusion that a curve $x(t) := (r_0(t), r_1(t))$ is a solution to the equations of motion if and only if $\frac{\partial L''}{\partial r_i} - \frac{d}{dt} \frac{\partial L''}{\partial \dot{r}_i} = 0$ is satisfied at each point on the curve for $i = 0, 1$. For $i = 1$, we again find that $\ddot{r}_1 = -\frac{1}{2r_1^2}$, so we get the same dynamics for the evolution of the relative distances as in the Newtonian and Leibnizean cases. But for $i = 0$, our condition on curves is empty, since L'' does not depend on either r_0 or \dot{r}_0 . It follows that a curve $x(t) := (r_0(t), r_1(t))$ counts as a solution to our equations of motion if $r_1(t)$ describes a motion permitted by the Newtonian or Leibnizean theory and r_0 is any (continuous and appropriately differentiable) function whatsoever.

Let us contrast the structure of these three theories.

Symmetries. The group of variational symmetries of the Newtonian theory is three-dimensional, consisting of Galilean boosts and spatial and temporal translations. The group of variational symmetries of the Leibnizean theory is one-dimensional, consisting of time translations. But the variational symmetry group of the semi-Leibnizean theory is infinite-dimensional: in addition to temporal translations, it includes time-dependent spatial translations of the centre of mass as a group of gauge symmetries. If $r_0(t)$ and $r_1(t)$ are continuous functions, then $x(t) = (r_0(t), r_1(t))$ is a kinematical possibility. Let $\Lambda(t)$ be any other continuous function from \mathbb{R} to itself. Then $x'(t) := (r_0(t) + \Lambda(t), r_1(t))$ is also a kinematical possibility and $L''(x(t)) = L''(x'(t))$ for all t (since L'' doesn't care at all about r_0). That is, the map $\Phi_\Lambda : (r_0(t), r_1(t)) \mapsto (r_0(t) + \Lambda(t), r_1(t))$ from the space of kinematical possibilities to itself preserves the Lagrangian, and hence is a variational symmetry. Indeed, for each such Λ we get a distinct variational symmetry of L'' . So the space of continuous $\Lambda : \mathbb{R} \rightarrow \mathbb{R}$ parameterizes a group of gauge symmetries of L'' .

Gauge Symmetries and the Initial Value Problem. We can exploit these symmetries to show how drastically ill-posed the initial value problem for the semi-Leibnizean theory is. Suppose that at $t = 0$ we fix values for $r_0, r_1, \dot{r}_0,$ and \dot{r}_1 . Let $x(t) = (r_0(t), r_1(t))$ be a solution satisfying those initial data. Now select $\Lambda : \mathbb{R} \rightarrow \mathbb{R}$ such that $\Lambda(0) = 0$ and $\dot{\Lambda}(0) = 0$. Since Φ_Λ is a Lagrangian symmetry, $\Phi_\Lambda(x(t)) = (r_0(t) + \Lambda(t), r_1(t))$ is also a solution — which, of course, satisfies the specified initial data at time $t = 0$. In this way, we can construct an infinite-dimensional family of solutions for each specified set of initial data.

Structure of the Space of Solutions. The spaces of solutions for the Newtonian and Leibnizean theories are symplectic spaces, of dimension four and two, respectively. As we have seen the space of solutions of the semi-Leibnizean theory is infinite-dimensional. And the form that L'' induces on this space is degenerate — the space is not symplectic. The associated gauge orbits have the following structure: $x(t) = (r_0(t), r_1(t))$ and $x'(t) = (r'_0(t), r'_1(t))$ lie in the same gauge orbit if and only if $r_1(t) = r'_1(t)$ for all t (i.e., solutions lie in the same gauge orbit if and only if they agree about the relative distances between the particles — what they say about the motion of the centre of mass is irrelevant).

Hamiltonian Picture. Writing $p_i = \dot{r}_i$, we find that the spaces of initial data for our theories are as follows.

1. For the Newtonian theory, the space of initial data is $T^*\mathbb{R}^2 = \{(r_0, r_1, p_0, p_1) : r_i, p_i \in \mathbb{R}\}$ carrying its canonical symplectic structure $\omega = \sum_{i=0,1} dr_i \wedge dp_i$. The Hamiltonian is $H(r_0, r_1, p_0, p_1) = \frac{1}{2}(p_0^2 + p_1^2) - \frac{1}{2r_1}$. The equations of motion are the usual deterministic Newtonian equations.
2. For the Leibnizean theory, the space of initial data is $T^*\mathbb{R} = \{(r_1, p_1) : r_1, p_1 \in \mathbb{R}\}$ carrying its canonical symplectic structure $\omega = dr_1 \wedge dp_1$. The Hamiltonian is $H(r_1, p_1) := \frac{1}{2}p_1^2 - \frac{1}{2r_1}$. The equations of motion are the usual deterministic Leibnizean equations.
3. Recall that in constructing the Hamiltonian system corresponding to a given Lagrangian system, we must first construct the momentum variables corresponding to the position variables of the Lagrangian system. The semi-Leibnizean theory has two position variables, r_0 and r_1 . Our recipe tells us that the corresponding momentum variables are $p_i := \frac{\partial L''}{\partial \dot{r}_i}$, for L'' the semi-Leibnizean Lagrangian. As usual, $p_1 := \dot{r}_1$. But because L'' is independent of \dot{r}_0 , we find that $p_0 \equiv 0$. It follows that the space of initial data for this theory is the space $\Gamma = \{(r_0, r_1, p_1) : r_0, r_1, p_1 \in \mathbb{R}\}$ that arises when we restrict attention to those states in the space of initial data for the Newtonian theory in which $p_0 = 0$; restricting the symplectic structure of the Newtonian theory to Γ yields a presymplectic structure (the vectors pointing in

the r_0 direction are the null vectors). The gauge orbits have the following structure: $x = (r_0, r_1, p_1)$ and $x' = (r'_0, r'_1, p'_1)$ lie in the same gauge orbit if and only if $r_1 = r'_1$ and $p_1 = p'_1$. The Hamiltonian for this theory is $H''(r_0, r_1, p_1) := \frac{1}{2}p_1^2 - \frac{1}{2r_1}$, which determines the usual Newtonian/Leibnizean behaviour for r_1 and p_1 while leaving the evolution of r_0 wholly unconstrained. That is, if $x(t)$ and $x'(t)$ are curves in the space of initial data corresponding to solutions of this Hamiltonian problem, then one finds that in general $x(t) \neq x'(t)$ for $t \neq 0$, but $[x(t)] = [x'(t)]$ for all t . Note that each such curve $x(t)$ corresponds to a point in the space of solutions, and that the condition $[x(t)] = [x'(t)]$ for all t just says that for the points in the space of solutions that correspond to the curves x and x' , themselves lie in the same gauge orbit.

Reduction. As one would expect, the reduced space of initial data of the semi-Leibnizean theory is isomorphic to the space of initial data of the Leibnizean theory, and the reduced space of solutions of semi-Leibnizean theory is isomorphic to the space of solutions of the Leibnizean theory — in both cases, this is because identifying points in the relevant gauge orbits amounts to dropping r_0 as a dynamical variable. So reduction implements our physical intuition that r_0 is an extraneous variable that ought to be excised and eliminates the pathologies of the semi-Leibnizean theory. Furthermore, the reduced space of initial data (reduced space of solutions) inherits from the original theory a Hamiltonian (Lagrangian) that is really that of the Leibnizean theory — so these reduced spaces carry dynamical theories with the correct dynamics and symmetry groups.

Of course, this is a toy example — one of the simplest possible. And it has been set up here so that it is clear from the beginning that the variables of the semi-Leibnizean theory can be segregated into the physically relevant r_1 , which plays a role in the Lagrangian and whose dynamics is deterministic, and the physically otiose r_0 , which plays no role in the Lagrangian, and whose evolution is completely unconstrained by the dynamics. So it has been clear from the beginning that r_0 ought to be excised from the theory — there has been no temptation to keep it on board and to conclude that we have an indeterministic theory on our hands.

But note that if we had stuck with our original Newtonian variables, q_1 and q_2 (with $q_1 < q_2$), and had written $L'' := \frac{1}{2}(\dot{q}_2 - \dot{q}_1)^2 - \frac{1}{2(q_2 - q_1)}$ then things would not have been quite so clear: the equations of evolution for q_1 and q_2 would have mixed together physically relevant information and physically otiose information and it would have taken a little bit of work to see what was going on.

When we are faced with Lagrangian theories admitting groups of local symmetries, we know (unless they exhibit the sort of pathological behaviour we saw in example 36 above) that there is some way of separating the variables into the physically relevant and the physically otiose (it is easiest to see this on the Hamiltonian side). But it is not always easy to find such a separation. This is one of

several reasons why we end up working with such theories rather than with the more attractive reduced theories that stand behind them.

EXAMPLE 38 (Maxwell’s Theory). We consider the electromagnetic field. Let V be Minkowski spacetime, and fix an inertial frame and an associated set of coordinates (x_0, x_1, x_2, x_3) . We choose as the target space for our field $W = \mathbb{R}^4$. So the kinematically possible fields are of the form $A : V \rightarrow \mathbb{R}^4$ (subject to some unspecified differentiability and boundary conditions). A is the usual four-potential.

We define $F_{\nu\mu} := \frac{\partial A_\nu}{\partial x_\mu} - \frac{\partial A_\mu}{\partial x_\nu}$ ($\nu, \mu = 0, \dots, 3$). So a kinematically possible field $A(x)$ determines a matrix-valued function, F . We label the component functions making up F according to the following scheme, thus identifying components of the F with components of the electric and magnetic fields, $\mathbf{E}(x) = (E^1(x), E^2(x), E^3(x))$ and $\mathbf{B}(x) = (B^1(x), B^2(x), B^3(x))$:

$$F_{\mu\nu}(x) = \begin{vmatrix} 0 & -E^1(x) & -E^2(x) & -E^3(x) \\ E^1(x) & 0 & B^3(x) & -B^2(x) \\ E^2(x) & -B^3(x) & 0 & B^1(x) \\ E^3(x) & B^2(x) & -B^1(x) & 0 \end{vmatrix}$$

We take as the Lagrangian for our theory $L := -\frac{1}{2} \left(|\mathbf{B}(x)|^2 + |\mathbf{E}(x)|^2 \right)$. Writing $A(x) = (A_0(x), A_1(x), A_2(x), A_3(x))$ and $\mathbf{A}(x) := (A_1(x), A_2(x), A_3(x))$, we find that the equations of motion for our Lagrangian are:

$$\begin{aligned} \nabla^2 A_0 + \frac{\partial}{\partial x_0} (\nabla \cdot \mathbf{A}) &= 0 \\ \nabla^2 \mathbf{A} - \frac{\partial^2 \mathbf{A}}{\partial x_0^2} &= 0 \end{aligned}$$

(here $\nabla := (\frac{\partial}{\partial x_1}, \frac{\partial}{\partial x_2}, \frac{\partial}{\partial x_3})$ is the ordinary three-dimensional gradient operator). These equations are equivalent to the usual vacuum Maxwell equations for the electric and magnetic fields: $\dot{\mathbf{B}} = -\nabla \times \mathbf{E}$, $\nabla \cdot \mathbf{B} = 0$, $\dot{\mathbf{E}} = \nabla \times \mathbf{B}$, and $\nabla \cdot \mathbf{E} = 0$.

Let $\Lambda : V \rightarrow \mathbb{R}$ be a continuous function (appropriately differentiable and satisfying appropriate boundary conditions). Then the map $\Phi_\Lambda : A \mapsto A' := A + d\Lambda$ is a map from the space of kinematically possible fields to itself. If one calculates the matrices F' and F corresponding to A and A' , one finds $F' = F$. So \mathbf{E} and \mathbf{B} are invariant under our gauge transformation $A \mapsto A'$. It follows that $L(\Phi_\Lambda(A)) - L(A) = 0$, so Φ_Λ is a Lagrangian symmetry — in particular, A' is a solution if and only if A is. Since Λ was arbitrary, and since Λ and Λ' lead to distinct symmetries so long as $d\Lambda \neq d\Lambda'$, we have in fact found a huge family of symmetries of our theory. Indeed, the Φ_Λ form a group of gauge symmetries of our theory in the official sense introduced above.

Of course, it follows that the initial value problem for A is ill-posed: let $A(x)$ be a solution for initial data posed on the instant $x_0 = 0$ and let Λ be a nonconstant function that vanishes on a neighbourhood of the hypersurface $x_0 = 0$; then A and $A' = A + d\Lambda$ are solutions that agree on $x_0 = 0$ but do not agree globally.

And, of course, the form that our Lagrangian induces on the space of solutions is degenerate. The corresponding gauge orbits have the following form: solutions A and A' belong to the same gauge orbit if and only if there is a $\Lambda : V \rightarrow \mathbb{R}$ such that $A' = A + d\Lambda$. An equivalent condition is that A and A' lie in the same gauge orbit if and only if they lead to the same \mathbf{E} and \mathbf{B} — which is just to say that the reduced space of solutions is the space of solutions to the field equations for \mathbf{E} and \mathbf{B} (remember, we are working in a fixed coordinate system, so these are well-defined). This reduced space is a symplectic manifold.

We can construct the Hamiltonian theory corresponding to our Lagrangian theory (our chosen inertial coordinates give us a slicing). For convenience, we take the configuration variables for our Lagrangian theory to be $A_0(x)$ and $\mathbf{A}(x)$. Let \mathcal{Q} be the space of possible (A_0, \mathbf{A}) and $T^*\mathcal{Q}$ be the corresponding cotangent bundle, carrying its canonical symplectic structure. A point in $T^*\mathcal{Q}$ consists of a quadruple $(A_0(x), \mathbf{A}(x), \pi_0(x), \boldsymbol{\pi}(x))$ of fields on spacetime, with A_0 and π_0 taking values in \mathbb{R} and \mathbf{A} and $\boldsymbol{\pi}$ taking values in \mathbb{R}^3 . Our usual procedure tells us that the momentum π_0 corresponding to A_0 is identically zero (our Lagrangian does not depend on \dot{A}_0); the momentum $\boldsymbol{\pi}$ corresponding to \mathbf{A} is $\boldsymbol{\pi}(x) = -\mathbf{E}(x)$. So the space Γ of initial data for our theory is the subspace $T^*\mathcal{Q}$ of points of the form $(A_0(x), \mathbf{A}(x), 0, \boldsymbol{\pi}(x))$ — so we can take points in Γ to be triples of the form $(A_0, \mathbf{A}, \boldsymbol{\pi})$. The presymplectic form that Γ inherits from its embedding in $T^*\mathcal{Q}$ yields gauge orbits of the following form: $(A_0, \mathbf{A}, \boldsymbol{\pi})$ and $(A'_0, \mathbf{A}', \boldsymbol{\pi}')$ belong to the same gauge orbit if and only if $\boldsymbol{\pi} = \boldsymbol{\pi}'$ and $\nabla \times \mathbf{A} = \nabla \times \mathbf{A}'$. Since $\boldsymbol{\pi} = -\mathbf{E}$ and $\mathbf{B} = \nabla \times \mathbf{A}$, this tells us that two points in the space of initial data lie in the same gauge orbit if and only if they correspond to the same electric and magnetic fields. It follows that the reduced space of initial data is just the space of instantaneous states of the electric and magnetic fields. We again find that the space of reduced space of initial data is symplectically isomorphic to the reduced space of solutions.

In present case, as in the semi-Leibnizean example above, we can view the given Lagrangian theory as containing surplus unphysical variables, whose evolution is undetermined by the dynamics, alongside physically sensible variables whose evolution is fully determined by the dynamics. In the present case, however, it is a bit harder to make this division explicit: clearly the good variables are the electric and magnetic fields and the bad ones are those that encode additional information in A — all we care about is which gauge orbit A lies in, so a specification of A gives us surplus information. Reduction allows us to avoid ever mentioning this sort of surplus information.

We can formulate a Hamiltonian version of Maxwell's theory in the setting of the reduced space of initial data: points in this space specify the values of the electric and magnetic fields at points of space at a given time; this space is symplectic; and it is possible to find a Hamiltonian on this space that drives the dynamical evolution encoded in Maxwell's equations for \mathbf{E} and \mathbf{B} .¹²⁰

It is natural to wonder whether the reduced space of solutions also supports a

¹²⁰See, e.g., [Marsden and Weinstein, 1982].

Lagrangian version of Maxwell's theory. That is, is there a Lagrangian in terms of \mathbf{E} , \mathbf{B} , and their derivatives whose variational problem has as its Euler–Lagrange equations the Maxwell's equations for \mathbf{E} and \mathbf{B} ?

At first sight, it might seem that we could just employ our original Lagrangian,

$$L := -\frac{1}{2} \left(|\mathbf{B}|^2 + |\mathbf{E}|^2 \right),$$

for this purpose. But this leads to the wrong equations of motion. And there is reason to worry that \mathbf{E} and \mathbf{B} are ill-suited to the variational approach, since their six components are not independent — they can be derived from the three-component vector potential — and so cannot be varied independently.¹²¹ So it is not obvious that the reduced space of solutions does support a Lagrangian version of Maxwell's theory.

Whether or not this problem is insuperable in the case of Maxwell's theory in Minkowski spacetime, other problems lie ahead. Suppose that we construct our spacetime, V , by rolling up one of the spatial dimensions of Minkowski spacetime: V is locally Minkowskian but has the global structure of $\mathbb{R}^3 \times S^1$. This makes a surprising difference to our theory. It is still true that the gauge orbits in the space of solutions are of the following form: $[A] := \{A + d\Lambda\}$ for all appropriate Λ . And it is still true that specifying a gauge orbit $[A]$ determines the behaviour of the electric and magnetic fields on spacetime. But it is no longer quite true that we can go in the other direction: in order to specify a gauge orbit $[A]$, one has to specify in addition to \mathbf{E} and \mathbf{B} also a single complex number, which we will call the *holonomy*. Intuitively, the holonomy measures the phase change that results when an electron is transported along a given loop that wraps once around the closed dimension of space. Thus a point in the reduced space of solutions can be viewed as consisting of a specification of \mathbf{E} and \mathbf{B} plus the holonomy. This extra number ruins everything: for while \mathbf{E} and \mathbf{B} are appropriately local objects, assigning a property to each point of spacetime, the holonomy is a nonlocal item. This becomes even more clear if we look for a way of describing the reduced space of solutions that does not have the strange feature of including two very different sorts of variable: the best way to proceed appears to be to describe a point in the reduced space of solutions as an certain sort of (highly constrained) assignment of a complex number to each closed curve in spacetime. So in such a topologically nontrivial spacetime, in order to specify a gauge orbit $[A]$ we need to specify nonlocal information. The present framework requires that a Lagrangian field theory involve an assignment of a property of each point of spacetime, and so cannot accommodate this example.¹²²

REMARK 39 (Lagrangians and the Reduced Space of Solutions). In the very simple particle theory considered in example 37 we saw a case in which the reduced

¹²¹See [Goldstein, 1953, 366] for this point. See [Sudbery, 1986] for a way around this worry — which, however, requires a slight generalization of the present notion of a Lagrangian theory.

¹²²That is, we seem to be talking about properties that require something bigger than a point to be instantiated, in violation of Humean supervenience (see fn. 45 above).

space of solutions of a theory admitting gauge symmetries inherited from the original theory a Lagrangian that encoded the gauge-invariant aspects of the original dynamics. But in the more interesting case of Maxwell's theory, considered in example 38, it seems less likely that there is any sense in which the reduced space of solutions arises directly from a local Lagrangian, without passing through a formulation admitting gauge symmetries. And this seems very unlikely indeed if we choose our spacetime to be topologically nontrivial, because in this case the Maxwell field appears to involve a non-local degrees of freedom.

Note that things become even worse in non-Abelian Yang–Mills theories. In these theories, the space of fields is the space of connection one-forms on a suitable principal bundle $P \rightarrow V$ over spacetime, the Lagrangian is a direct generalization of the Lagrangian of Maxwell's theory, and the group of gauge symmetries is the group of vertical automorphisms of P . The reduced space of solutions is the space of connections modulo vertical automorphisms of P . Even when V is Minkowski spacetime, the best parameterization of the reduced space of solutions would appear to be one that deals with holonomies around closed curves in spacetime.¹²³ So it would again appear difficult (perhaps impossible) to capture this reduced space of solutions via the variational problem of a local Lagrangian.¹²⁴ Indeed, it seems plausible the prevalence of gauge freedom in physical theories is grounded in the fact that by including nonphysical variables one is sometimes able to cast an intrinsically nonlocal theory in to a local form.¹²⁵

6.3 Time-Dependent Systems

Let us assume that our spacetime, V , admits a slicing, and that our equations of motion, Δ , are second-order and exhibit good existence and uniqueness properties.¹²⁶ But we now assume that our Lagrangian L is *time-dependent*, in the sense that it does not admit a dynamical time translation group, τ , arising from a time translation group, $\bar{\tau}$, on V .

The time-dependent Lagrangian theories that arise in physics fall under the two following cases.

¹²³There is, however, considerable controversy among philosophers regarding the best interpretation of classical non-Abelian Yang–Mills theories. See [Healey, Unpublished], [Maudlin, Unpublished], and [Belot, 2003, §12].

¹²⁴Under a usage distinct from the present one, any Hamiltonian theory on a velocity phase space (i.e., a tangent bundle) counts as a Lagrangian theory; see, e.g., [Abraham *et al.*, 1988, Chapter 8]. Under this alternative use, Lagrangians are not required to be local and a variational principle plays no necessary role. It may well be that there are treatments of theories that are Lagrangian in this sense, but not in the sense that I am concerned with in the text.

¹²⁵On this point, see, e.g., [Belot, 2003, §13]. For further speculation about the importance of gauge freedom, see [Redhead, 2003].

¹²⁶Recall from section 5.2 above that a slicing of a spacetime is a decomposition into space and time; not every slicing satisfies the stronger condition that this decomposition meshes with a time translation group on V . Only spacetimes with geometries strong enough to determine a family of instants and a family of possible point-particle worldlines admit slicings.

Case (A): V admits a time translation group $\bar{\tau}$, but this group does not correspond to a symmetry of the equations of motion. Example: A system of particles in Newtonian spacetime, subject to forces arising from a time-dependent potential.

Case (B): V does not admit a time translation group. For example, let (V, g) be a curved general relativistic spacetime without temporal symmetries and take the Klein–Gordon equation for a scalar field on (V, g) , $\nabla^a \nabla_a \Phi - m^2 \Phi = 0$, as the equation of motion (note that the metric on V plays a role in defining the derivative operators); the corresponding Lagrangian is $L = \frac{1}{2} \sqrt{-g} \nabla_a \nabla^a \Phi + m^2 \Phi^2$.

We will also need that notion of a time-dependent Hamiltonian system.

DEFINITION 40 (Time-Dependent Hamiltonian Systems). A *time-dependent Hamiltonian system* (M, ω, h) consists of a symplectic manifold (M, ω) , called the *phase space*, together with a smooth function $h : \mathbb{R} \times M \rightarrow \mathbb{R}$, called the *Hamiltonian*. We often write $h(t)$ for $h(t, \cdot) : M \rightarrow \mathbb{R}$.

Ordinary Hamiltonian systems (see definition 8 above) are special cases of time-dependent Hamiltonian systems in which $h(t)$ is the same function on M for each value of t ; we will also call such systems *time-independent Hamiltonian systems*. In a time-independent system, the dynamical trajectories could be thought of as curves in the phase space, parameterized up to a choice of origin, with exactly one such curve passing through each point of the space. In the time-dependent case, the situation is more complicated. For each value of t , we can solve $\omega(X_{h(t)}, \cdot) = dh(t)$ for the vector field $X_{h(t)}$ generated by $h(t)$. We can then declare that a curve $\gamma : \mathbb{R} \rightarrow \mathcal{I}$ is a dynamical trajectory of (\mathcal{I}, ω, h) if for each $t \in \mathbb{R}$, $\dot{\gamma}(t) = X_{h(t)}(\gamma(t))$ — that is, for each t , the tangent vector to γ at $x = \gamma(t)$ is given by the value of the vector field $X_{h(t)}$ at x . Notice that while in the case of a time-independent Hamiltonian system, there is a single dynamical trajectory through each point of the phase space, in the present case there will in general be many such trajectories through each point (since which states come immediately after $x \in \mathcal{I}$ depends on the tangent to the dynamical trajectory through x ; and in the time-dependent case, this tangent will vary as we consider posing initial data x at different possible instants).

Given the set of assumptions that we have in play, we expect to find the following when we investigate a time-dependent Lagrangian theory.

Lagrangian Picture. One can apply the usual variational procedure to pass from a Lagrangian to a set of equations of motion. We can also follow the usual procedure in order to equip the corresponding space of solutions, \mathcal{S} , with a two-form, Ω — and, as usual, one presumes that for the sort of examples that arise in physics, uniqueness of solutions to the equations of motion implies that Ω is symplectic.¹²⁷ Note, however, that in time-dependent the-

¹²⁷For a discussion of the construction of (\mathcal{S}, Ω) in the time-dependent case, see [Woodhouse, 1991, §2.4].

ories of the types under consideration, \mathcal{S} does *not* carry a one-parameter group implementing time translations: in theories falling under Case (A) above, such a group acts on the space of kinematically possible fields, but (in general) maps solutions to non-solutions; in theories falling under Case (B), there is no available notion of time translation. We can as usual use the stress-energy tensor of the field to define the energy of the field along any given instant — but the result is no longer independent of the instant chosen.

Hamiltonian Picture. A choice of slicing for our spacetime V leads to a Hamiltonian picture which is in many ways similar to that which emerges in the time-independent case. Let S be a manifold homeomorphic to an arbitrary instant $\Sigma \subset V$ (and with the geometry, if any, shared by such instants) and let σ be a slicing of V employing S as an abstract instant: it is helpful to think of the choice of σ as the choice of a preferred family of observers equipped with a notion of simultaneity. Then we can set about constructing a Hamiltonian version of our theory, following in so far as possible the recipe from the time-independent case.

1. Given an instant Σ in our slicing and a solution Φ , we define: ϕ , the restriction of the field to Σ ; $\dot{\phi}$ the time rate of change of Φ along Σ relative to the observers and clocks that define σ ; and $\pi := \frac{\partial L}{\partial \dot{\phi}}$, the field momentum along Σ relative to the slicing σ .
2. Given a solution Φ and instant Σ in our slicing, we use σ to pull back to S the initial data (ϕ, π) induced by Φ on Σ , and henceforth think of ϕ and π as functions defined on S , when convenient.
3. Let \mathcal{Q} be the space of all $\phi : S \rightarrow W$ that arise in this way; then $T^*\mathcal{Q}$ is the space of all pairs (ϕ, π) that arise in this way. This is our space of initial data, \mathcal{I} . It carries a canonical symplectic form, ω (see example 7 above).
4. The construction of the Hamiltonian is the first stage at which we run into any novelty.¹²⁸ Let Σ_t be an instant in our slicing, and define $h(t) : \mathcal{I} \rightarrow \mathbb{R}$ (t fixed, for now) as $h(t)(\phi, \pi) := \int_{\Sigma_t} \pi \dot{\phi} - L(\Phi) dx$, where Φ is the solution that induces (ϕ, π) on Σ_t , and $\dot{\phi}$ is the field velocity that Φ induces on Σ_t . In general, this construction yields a different real-valued function on \mathcal{I} for each value of t . One expects that $h(t)(\phi, \pi)$ gives the total instantaneous energy when initial data (ϕ, π) are posed on instant Σ_t . But imposing the same initial data at distinct times in general leads to states with different total energies (since, roughly speaking, we are dealing with systems subject to time-dependent potentials).
5. Now considering t as a variable, we see that we have defined a smooth $h : \mathbb{R} \times \mathcal{I} \rightarrow \mathbb{R}$. So (\mathcal{I}, ω, h) is a time-dependent Hamiltonian system

¹²⁸For this construction see, e.g., [Kay, 1980, §1].

in the sense of definition 40 above. The resulting dynamics can be thought of as follows. Suppose that we are interested in the dynamics that results when we pose our initial data on a fixed instant Σ_{t_0} in our slicing. Then, for each $s \in \mathbb{R}$ we can ask what state $x \in \mathcal{I}$, posed on Σ_{t_0} , evolves into after s units of time; we call the result $g_s^{t_0}(x)$. This gives us a map $g_s^{t_0} : \mathcal{I} \rightarrow \mathcal{I}$ for each s ; and the set $\{g_s^{t_0}\}_{s \in \mathbb{R}}$ forms a one-parameter group; each of $g_s^{t_0}$ is a symplectic automorphism of \mathcal{I} but does not leave h invariant. So here we have the dynamics implemented by symmetries of (\mathcal{I}, ω) that are not symmetries of (\mathcal{I}, ω, h) . Letting t_0 vary gives us a one-parameter family of such one-parameter dynamics-implementing groups.

Relation between the Pictures. As in the time-independent setting, for each Σ in our slicing, we can define the map $T_\Sigma : \mathcal{S} \rightarrow \mathcal{I}$ that sends a solution to the initial data it induces on Σ . Because we are assuming global existence and uniqueness for solutions to our equations of motion, each such map is a bijection. Furthermore, as in the time-independent case, each such T_Σ is in fact a symplectic isomorphism between \mathcal{S} and \mathcal{I} . We can use these maps to show that the time-dependent Hamiltonian system constructed above encodes the correct dynamics for our equations of motion: let Φ be a solution and let x_0 be the initial data induced by Φ on the instant Σ_0 , and let $x_0(t)$ be the corresponding dynamical trajectory in the space of initial data; then for each $t \in \mathbb{R}$, $x_0(t)$ is the initial data that Φ induces on Σ_t .

In the time-independent case, we also found that the maps T_Σ intertwined the actions of the group implementing time translation on the space of solutions and time evolution on the space of initial data. In the present case, we have, so far, nothing corresponding to time translation on the space of solutions, while on the space of initial data, we have a whole family of notions of time evolution (indexed by a choice of instant upon which initial data are to be posed). Now note that for each instant Σ_t in our slicing and each $s \in \mathbb{R}$ we can define $\hat{g}_s^t := T_{\Sigma_t}^{-1} \circ g_s^t \circ T_{\Sigma_t}$; the family $\{\hat{g}_s^t\}_{s \in \mathbb{R}}$ is a one-parameter group of symplectic automorphisms of (\mathcal{S}, Ω) which is *not* a group of variational symmetries of our Lagrangian. The result of applying \hat{g}_s^t to a solution Φ is the solution that would result if the initial data induced by Φ on Σ_t had been posed instead on the instant Σ_{t-s} .¹²⁹

Time. In the present context, time translation may or may not be a symmetry of our spacetime. But even if it is, there is no corresponding symmetry of the dynamics. And so our picture is hobbled — we do not get nice actions of the real numbers on the space of solutions and on the space of initial data that implement time translation and time evolution. On the space of initial

¹²⁹Of course, in the time-independent case, this reduces to time translation of solutions — so we can regard the transformations \hat{g}_s^t as generalizing the ordinary notion of time translation of solutions.

data, for each instant at which we might choose to pose initial data, we get a one-parameter group implementing time evolution — but this is not a symmetry of the Hamiltonian. On the space of solutions, we have no natural group action corresponding to time-translation. If we choose a slicing and an instant, then we can get an \mathbb{R} -action that gives us information not about time translation of solutions, but about what solution results if we take the initial data that a given solution induces on that given instant and re-pose it on another instant.

Change. Some physical quantities will be represented by functions on the space of initial data: for example, in a theory of two Newtonian particles subject to time-dependent external forces, relative distance between the particles will be encoded in a function on the space of initial data. But some quantities will be represented by one-parameter families of functions on the space of initial data: energy will be an example of such a quantity in any time-dependent system.¹³⁰ As we have done above with the Hamiltonian, let us use the symbol $f(t)$ to denote such a one-parameter family — we can think of an ordinary function as being a degenerate case, where $f(t)$ is the same function on \mathcal{I} for each $t \in \mathbb{R}$. Let $x(t)$ be a dynamical trajectory in \mathcal{I} . Then $x(t)$ represents the quantity modelled by $f(t)$ as changing if and only if $\exists t_1, t_2 \in \mathbb{R}$ such that $f(t_1)(x(t_1)) \neq f(t_2)(x(t_2))$.

On the space of solutions, we expect that, once we have chosen a slicing, each quantity of interest will be represented as usual by a one-parameter family of functions — as usual, we denote such a family of functions on \mathcal{S} by $\{f_t\}$. Suppose that a quantity of interest is represented by $f(t)$ on the space of initial data, and let Σ_{t_0} be an instant in our slicing. Then we define $f_{t_0} := f(t_0) \circ T_{\Sigma_{t_0}} : \mathcal{S} \rightarrow \mathbb{R}$. Carrying this out for each $t \in \mathbb{R}$ gives us our desired $\{f_t\}$. As usual, we view a solution $\Phi \in \mathcal{S}$ as representing our quantity as changing if $\exists t_1, t_2 \in \mathbb{R}$ such that $f_{t_1}(\Phi) \neq f_{t_2}(\Phi)$.

REMARK 41 (Artificially Time-Dependent Theories). If we have a time-independent Lagrangian theory but perversely choose a slicing that is not adapted to our notion of time translation, then the result of following the above procedure would be a time-dependent Hamiltonian system.

REMARK 42 (Quantization of Time-Dependent Systems). There is no special difficulty in quantizing a time-dependent Hamiltonian treatment of a system with finitely many degrees of freedom. But it is not in general possible to construct a well-behaved quantum Hamiltonian for a time-dependent field theory.¹³¹ For

¹³⁰It is not hard to find other examples. In a field theory set in a nonstationary spacetime, the abstract instant will not carry a Riemannian metric (since the instants $\Sigma \subset V$ do not share a Riemannian geometry). In this case, we find that an initial data set that represents the field as having two sharp peaks will correspond to instantaneous states in which the peaks are different distances apart, depending on the instant Σ_t in the slicing upon which the initial data are posed. So in this sort of example, even relative distance is represented by a family of functions on the space of initial data.

¹³¹See, e.g., [Kay, 1980, §2.1].

this reason, the standard construction of free quantum field theories on curved spacetimes take as their starting point the space of solutions rather than the space of initial data.¹³²

7 THE PROBLEM OF TIME IN GENERAL RELATIVITY

General relativity differs from the theories considered above in being generally covariant. It is widely accepted that this leads to certain characteristic technical and conceptual problems, grouped together under the rubric *the problem of time*. This section forms an extended commentary on the problem of time in general relativity. The first subsection below is devoted to a discussion of the general covariance of general relativity and some of its direct consequences. The following subsection contains a discussion of the problem of time itself — essentially that change cannot be represented in the theory in the way familiar from the discussion of sections 5 and 6 above. The final subsection discusses a strategy for finding time and change in general relativity (this discussion is intended by way of further clarification of the problem of time, rather than as a suggested resolution).

It is important to emphasize that while the present discussion focuses on general relativity, the problems under discussion arise whenever one has a theory that is generally covariant in an appropriate sense.

7.1 *The General Covariance of General Relativity*

Let V be a spacetime manifold, with or without geometrical structure. Recall that a C^k ($0 < k \leq \infty$) diffeomorphism $d : V \rightarrow V$ is a C^k bijection with C^k inverse.¹³³ Leaving the degree of differentiability unspecified, we will denote by $\mathcal{D}(V)$ the group of diffeomorphisms from V to itself.¹³⁴

¹³²See [Wald, 1994, Chapter 4].

¹³³A diffeomorphism $d : V \rightarrow V$ is called *small* if it is homotopic to the identity, otherwise it is *large*. For ease of exposition, I implicitly restrict attention to small diffeomorphisms below. I will often speak of the pullback of a tensor by a diffeomorphism. The most important case will be the pullback d^*g of a spacetime metric g by a diffeomorphism d . Intuitively (V, d^*g) is the spacetime geometry that results if we lift g off V , then use d to permute the identities of points of V , then lay g back down. (V, g) and (V, d^*g) share a set of spacetime points and have isomorphic geometries; they differ only as to which points in V play which geometric roles — unless d is a symmetry of the metric, in which case they do not differ even about this.

¹³⁴Special care is required in dealing with groups of diffeomorphisms: on the one hand, the group of C^k diffeomorphisms from a compact manifold to itself has a nice differentiable structure — it is a Banach manifold — but is not a Banach Lie group because the operation of group multiplication is not smooth; on the other hand, the the group of C^∞ diffeomorphisms from a compact manifold to itself has a less desirable differentiable structure — it is a mere Fréchet manifold — but it is a Fréchet Lie group; see [Adams *et al.*, 1985] and [Milnor, 1984] for details. The situation is even worse for groups of diffeomorphisms from a noncompact manifold V to itself: it appears that one needs to presuppose some geometrical structure on V in order to give the group a differentiable structure; see [Cantor, 1979] and [Eichhorn, 1993]. See [Isenberg and Marsden, 1982] for tactics for circumventing these difficulties.

Roughly speaking, we want to say that a theory is generally covariant when it has $\mathcal{D}(V)$ as a symmetry group.¹³⁵ So for each of the several notions of symmetry of a theory, we have a corresponding notion of general covariance. Following [Earman, 2006], I will single out the following two as the most important for our purposes:

Weak General Covariance: $\mathcal{D}(V)$ is a group of symmetries of the equations Δ of the theory.

Strong General Covariance: $\mathcal{D}(V)$ is a group of gauge symmetries of the Lagrangian L of the theory.

Of course, Strong General Covariance implies Weak General Covariance (since every gauge symmetry is a variational symmetry, and therefore a symmetry of the equations of motion). But the converse is not true (a theory may be weakly generally covariant even if it does not admit a Lagrangian, and hence is not eligible to be strongly generally covariant).

We have been allowing V to carry a fixed geometrical background, encoded in some tensors that do not vary from solution to solution. We could have allowed V to carry further nongeometric solution-independent structure.¹³⁶ On the other hand, in any nontrivial theory the fields governed by the equations of motion will of course vary from solution to solution. So we have a distinction between theories in which V is equipped with nontrivial solution-independent structure and theories in which it is not.¹³⁷

Intuitively, a theory is weakly generally covariant if and only if its solutions carry no solution-independent tensors (or spinors, or ...) — for it is precisely when we have some fixed background tensors painted on V that the equations of motion can “care” about the distinction between a solution Φ and its pullback $d^*\Phi$ by a diffeomorphism $d: V \rightarrow V$.

Of course, general relativity is weakly generally covariant — indeed, in the important vacuum sector of general relativity the spacetime metric is the only basic quantity of the theory, and it is solution-dependent.¹³⁸

The question whether general relativity satisfies Strong General Covariance is a bit more subtle. Intuitively, it ought to: for at the formal level diffeomorphisms of V are variational symmetries of the Lagrangian of the theory, and the group

¹³⁵See [Norton, 1995] for the tangled history of the notion of general covariance.

¹³⁶For example, in studying the motion of charged matter in a strong external electromagnetic field, we might employ a theory in which the Maxwell field as well as the spacetime geometry was solution-independent and only the motion of the matter varied from solution to solution.

¹³⁷Note that the distinction made here between solution-independent and solution-dependent structures does not coincide with the Anderson–Friedman distinction between absolute and dynamical objects (see [Friedman, 1983, §II.2]): solution-independent objects are required to be the same from solution to solution while absolute objects are only required to be the same from solution to solution *up to diffeomorphism*.

¹³⁸In this regime, the Einstein Field equations just tell us that if metric g on V counts as a solution if and only if the Ricci curvature tensor of g vanishes. And clearly g is Ricci-flat if and only if d^*g is. So $\mathcal{D}(V)$ maps solutions to solutions.

of such diffeomorphisms is parameterized in a suitable sense by the set of vector fields on V . But, as we will soon see, this is a point at which technicalities about boundary conditions cannot be ignored.

But we can block out such technicalities by restricting attention to the subset $\mathcal{D}_c(V) \subset \mathcal{D}(V)$, consisting of compactly supported diffeomorphism from V to itself.¹³⁹ $\mathcal{D}_c(V)$ turns out to be a group of gauge symmetries of the Lagrangian of general relativity ($\mathcal{D}_c(V)$ is parameterized by the family of compactly supported vector fields on V). So the counterpart of condition (2) above goes through when $\mathcal{D}(V)$ is replaced by $\mathcal{D}_c(V)$.

In order to say more, and to approach the question of the significance of general covariance for questions about time, we turn below to two special cases: (i) general relativity in the spatially compact domain; (ii) general relativity in the domain in which asymptotic flatness is imposed at spacelike infinity. The first case is central to cosmology: by requiring space to be compact, one eliminates worries about boundary conditions at spatial infinity; this permits one to investigate universes packed with matter while maintaining control over technical issues. The second case is of more strictly mathematical and conceptual interest (the asymptotic boundary conditions of greatest physical interest impose asymptotic flatness at null infinity rather than spatial infinity; these allow one to investigate gravitational radiation). After discussing these cases, I briefly turn to the question whether every theory can be given a generally covariant formulation.

General Relativity as a Cosmological Theory

We restrict attention to vacuum general relativity in which the spacetime metric, g , is the only field. So we take as our space of kinematically possible fields the space of Lorentz signature metrics on some fixed n -dimensional spacetime manifold V .¹⁴⁰ The equation of motion for this theory is $R_{ab} - \frac{1}{2}Rg_{ab} = 0$, where R_{ab} is the Ricci curvature of g and R is the scalar curvature of g ; here and throughout we require the cosmological constant to vanish.

Recall that a subset $\Sigma \subset V$ is called a *Cauchy surface* of (V, g) if every inextendible timelike curve in (V, g) intersects Σ exactly once; it follows that a Cauchy surface is an $(n - 1)$ -dimensional spacelike submanifold of V . We call (V, g) *globally hyperbolic* if it possesses a Cauchy surface. If (V, g) is globally hyperbolic, then it can be foliated by Cauchy surfaces, and all of its Cauchy surfaces are homeomorphic to one another. Indeed, if (V, g) is globally hyperbolic, then V is homeomorphic to a manifold of the form $S \times \mathbb{R}$ for some $(n - 1)$ -dimensional manifold S , with all the Cauchy surfaces of (V, g) homeomorphic to S . For the purposes of this discussion of general relativity as a cosmological theory we restrict attention to solutions with compact and orientable Cauchy surfaces.¹⁴¹

¹³⁹That is, a diffeomorphism $d : V \rightarrow V$ is in $\mathcal{D}_c(V)$ if and only if there exists a compact set $U \subset V$ such that d acts as the identity on V/U .

¹⁴⁰So a kinematically possible field is a section of the bundle of symmetric bilinear forms of Lorentz signature over V .

¹⁴¹The restriction to globally hyperbolic solutions is not required for construction of a La-

We can proceed to construct Lagrangian and Hamiltonian versions of our theory.

Lagrangian Picture. The Lagrangian for general relativity is the *Einstein–Hilbert Lagrangian*, $L = \sqrt{-g}R$. The space of solutions is, of course, infinite-dimensional. Let us call a solution *well-behaved* if it admits a foliation by Cauchy surfaces with constant mean curvature.¹⁴² It is believed that the set of well-behaved solutions forms a large open subset of the full space of solutions; and it is known that within the space of well-behaved solutions the only singularities that occur are mild ones at metrics that admit Killing fields (these are vector fields that can be thought of as the infinitesimal generators of spacetime symmetries).¹⁴³ The group $\mathcal{D}(V)$ is a group of gauge symmetries of the Einstein–Hilbert Lagrangian: each one-parameter group of diffeomorphisms from V to itself is a group of variational symmetries of this Lagrangian, and the group $\mathcal{D}(V)$ can be parameterized by arbitrary vector fields on V .¹⁴⁴ So, in accord with the theory of gauge theories developed in section 6.2 above, we find that the space, \mathcal{S} , of well-behaved solutions carries a presymplectic form, Ω (henceforth I drop the qualifier and speak of \mathcal{S} as the space of solutions).¹⁴⁵ As usual, this presymplectic form induces a partition of the space of solutions by gauge orbits. Two metrics, g and g' , belong to the same gauge orbit if and only if there exists a diffeomorphism $d : V \rightarrow V$ such that $g' = d^*g$. Of course, the conserved quantities associated with one-parameter groups of diffeomorphism are trivial — each is the zero function on \mathcal{S} . Indeed, in this context, general relativity has no nontrivial Noether quantities — beyond diffeomorphisms, the only continuous, local symmetries of the laws are metric rescalings, which are not variational symmetries.¹⁴⁶

Reduced Space of Solutions. The space, \mathcal{S}' , of gauge orbits of the space of solutions of general relativity is a symplectic space with mild singularities at points corresponding to solutions with Killing fields.¹⁴⁷ Let us call a point $[g]$ in the reduced space of solutions a *geometry* — since distinct representatives of $[g]$ represent V as having the same spacetime geometry, but differ as to the distribution of geometrical roles to points of V . So far as I know, it makes no sense to speak of this reduced space as the space of solutions arising from

grangian version of general relativity, but is required for the Hamiltonian treatment and plays a role in some of the results cited below concerning the structure of the space of solutions. The requirement that the spatial topology be orientable is required for the Hamiltonian treatment.

¹⁴²Mean curvature will be defined below on p. 201, in the course of the discussion of the space of initial data.

¹⁴³For the structure of the space of well-behaved solutions, see [Isenberg and Marsden, 1982].

¹⁴⁴See [Crnković and Witten, 1987] and [Woodhouse, 1991, 143–146]; the latter provides an argument that non-compactly supported diffeomorphisms belong in the group of gauge symmetries of the Lagrangian.

¹⁴⁵See also [Frauendiener and Sparling, 1992] for a construction of the presymplectic form on the space of solutions which does not proceed via the Lagrangian formalism.

¹⁴⁶See [Torre and Anderson, 1996, esp. p. 489].

¹⁴⁷See [Isenberg and Marsden, 1982].

the variational problem for a local Lagrangian. Indeed, a geometry $[g]$ would not appear to assign any particular local property to any point $x \in V$.

Hamiltonian Picture. The construction of the corresponding Hamiltonian picture requires a bit of care.¹⁴⁸ We want to mimic as much of the procedure of section 5 above as we can, given that we do not have available a slicing (which requires that spacetime have a nontrivial solution-independent geometry). We proceed as follows.¹⁴⁹

1. *Construct the space of initial data.* Up until now, we have been able to proceed as follows: (i) choose a slicing σ of V and an instant $\Sigma \subset V$ in σ , then construct the space of possible instantaneous field configurations, \mathcal{Q} , by looking at all the $\phi : \Sigma \rightarrow W$ that arise by restricting solutions Φ to Σ ; (ii) construct the space of initial data $\mathcal{I} \subseteq T^*\mathcal{Q}$ by finding all pairs (ϕ, π) that are induced as initial data on Σ (where π is the instantaneous field momentum, defined via $\pi := \frac{\partial L}{\partial \dot{\phi}}$, with $\dot{\phi}$ is the time rate of change of the field according to the observers associated with the slicing σ). We found that \mathcal{I} was a proper subset of $T^*\mathcal{Q}$ whenever the Lagrangian L of the theory admitted a group of gauge symmetries. Without relying on a notion of slicing, we can construct a space of initial data via a procedure surprisingly close to the usual one.

If $\Sigma \subset V$ is a hypersurface and g is a solution to the Einstein field equations, then $q := g|_{\Sigma}$ is a symmetric covariant tensor of rank two. But in the present setting, the restriction of a solution to an arbitrary hypersurface is not a good candidate for an instantaneous configuration of the field: intuitively, since the gravitational field of general relativity is a spacetime geometry, an instantaneous configuration of this field should be a spatial geometry. But, of course, $q := g|_{\Sigma}$ is a Riemannian metric on Σ if and only if Σ is spacelike according to g . So $q = g|_{\Sigma}$ represents an instantaneous state of the field if and only if Σ is spacelike.¹⁵⁰ So it seems reasonable to take as the space of possible instantaneous field configurations, \mathcal{Q} , the space of Riemannian metrics q that arise by restricting each solution to the hypersurfaces that it renders spacelike.¹⁵¹

¹⁴⁸For the constructions that follow see [Wald, 1984, Appendix E.1] or [Beig, 1994].

¹⁴⁹The construction sketched here does not rely on the lapse and shift fields. Fixing the behaviour of these nonphysical fields allows one to pass from initial data on an abstract instant S to a solution on $S \times I$ for some (possibly small) interval I of real numbers. As such they allow one to establish a bijection between the space of initial data and a set of solutions of limited temporal extent. I avoid the lapse and shift here because I want to concentrate on global results and on physical fields. For a very helpful introduction to the lapse and shift formalism, see [Marsden *et al.*, 1972, §III].

¹⁵⁰In the spatially compact globally hyperbolic regime, a submanifold of (V, g) with the topology of a Cauchy surface for (V, g) is a Cauchy surface if and only if it is spacelike according to g ; see [Budic *et al.*, 1978, Theorem 1].

¹⁵¹The field q is taken to be defined on an abstract instant, S , diffeomorphic to the Cauchy surfaces $\Sigma \subset V$. In order to construct \mathcal{Q} we choose an instant Σ relative to (V, g) and a diffeo-

The definition of the instantaneous field momenta is more complicated. In the familiar case, the slicing σ plays an important role. But nothing like that is available in the present case: it is awkward to introduce a solution-independent notion of slicing in the context of general relativity, considered as a dynamical theory.¹⁵² There is, however, a way around this difficulty. Consider a solution g and a hypersurface $\Sigma \subset V$ that g represents as being spacelike. Relative to g we can choose a slicing of V in the usual sense (since relative to g we can single out the instants and possible worldlines of point-particles as submanifolds of V). We call such a slicing *Gaussian* for Σ if it corresponds to a set of freely falling observers whose clocks all read zero as they pass through Σ , and whose worldlines are all orthogonal to Σ . For sufficiently small t the hypersurfaces of constant t according to the Gaussian observers will be Cauchy surfaces carrying Riemannian metrics $q(t) := g|_{\Sigma_t}$. So given a Gaussian slicing for Σ , we can define $\dot{q}_{ab} := \frac{\partial q_{ab}(t)}{\partial t} |_{t=0}$, which is a symmetric covariant tensor of rank two on Σ . In fact, \dot{q}_{ab} is independent of the Gaussian slicing chosen, and can be viewed as telling us about the geometry of the embedding of Σ in (V, g) . We can take a similar view of the *extrinsic curvature* of Σ in (V, g) , $k_{ab} := \frac{1}{2} \dot{q}_{ab}$, and the *mean curvature* along Σ in (V, g) , $k := q^{ab} k_{ab}$. Now: relative to our Gaussian slicing, the tensor $\dot{q}_{ab}(0)$ represents the velocity of the gravitational field, in the sense that it encodes information about the time rate of change of the field; as usual, we can define the corresponding momentum as $\pi^{ab} := \frac{\partial L}{\partial \dot{q}}$. For the Einstein-Hilbert Lagrangian we have $\pi^{ab} = (\sqrt{q} k^{ab} - k q^{ab})$ (so the momentum is a symmetric contravariant tensor of rank two).¹⁵³ We take as our space of initial data the space $\mathcal{I} \subset T^*\mathcal{Q}$ of pairs (q, π) that arise as the field configuration and momentum induced by solutions on hypersurfaces they render spacelike (with q and π functions living on an abstract manifold S). As is to be expected in a theory with gauge symmetries, \mathcal{I} is a proper subspace of $T^*\mathcal{Q}$ and the restriction of the canonical symplectic form on $T^*\mathcal{Q}$ equips \mathcal{I} with a presymplectic form, ω . The gauge orbits of ω have the

morphism $d : S \rightarrow \Sigma$, and use d to pull back to S all of the q that arise as restrictions of Σ of solutions that render Σ spacelike. \mathcal{Q} is of course independent of the choice of Σ and d .

¹⁵²Suppose that σ is a slicing relative to a metric g on V . Then the restriction of g to the instants Σ_t in σ will be reasonable instantaneous field configurations, and so relative to σ the solution g ought to correspond to a curve in the space of initial data of the theory. But what happens if we look at another solution g' relative to σ ? In general, the result of restricting this new solution to Σ_t in σ will not be an instantaneous state of the field — so σ will not give us the means to associate with each spacetime solution g a trajectory in the space of initial data.

¹⁵³If we calculate \dot{q} relative to a non-Gaussian slicing of (V, g) , then we will in general get an answer quite different from that generated by a Gaussian slicing. But if we use this new notion of the field velocity in our definition of the field momentum, we find that our new observers agree with our original Gaussian observers about the value of the field momentum at each point of Σ . So, rather surprisingly, in general relativity the field momentum depends on the instant chosen, but not on a slicing.

following structure: initial data sets (q, π) and (q', π') belong to the same gauge orbit if and only if they arise as initial data for the same solution g .¹⁵⁴

2. *Construct a Hamiltonian.* Application of the usual rule for constructing a Hamiltonian given a Lagrangian leads to the Hamiltonian $h \equiv 0$.
3. *Construct dynamics.* Imposing the usual dynamical equation, according to which the dynamical trajectories are generated by the vector field(s) X_h solving $\omega(X_h, \cdot) = dh$, leads to the conclusion that dynamical trajectories are those curves generated by null vector fields. So a curve in \mathcal{I} is a dynamical trajectory if and only if it stays always in the same gauge orbit. This is, of course, physically useless — since normally we expect dynamical trajectories for a theory with gauge symmetries to encode physical information by passing from gauge orbit to gauge orbit. But in the present case, nothing else could have been hoped for. A non-zero Hamiltonian would have led to dynamical trajectories which passed from gauge orbit to gauge orbit — but this would have been physical nonsense (and worse than useless). For such dynamics would have carried us from an initial state that could be thought of as an instantaneous state for solution g to a later instantaneous state that could not be thought of as an instantaneous state for solution g . In doing so, it would have turned out to encode dynamical information very different from that encoded in Einstein's field equations.

Reduced Space of Initial Data. We can pass to, \mathcal{I}' , the space of reduced initial data: a point in this space consists of a gauge equivalence class of points in the space of initial data. Like the reduced space of solutions, the reduced space of initial data is a symplectic space with mild singularities.¹⁵⁵ Indeed, it is presumed that the two reduced spaces are canonically isomorphic as symplectic spaces, under the map that takes a gauge orbit of initial data to corresponding gauge orbit of solutions.¹⁵⁶ \mathcal{I}' inherits from \mathcal{I} the trivial Hamiltonian $h \equiv 0$; this induces the trivial dynamics on \mathcal{I}' according to which the dynamical trajectories are constant curves of the form $x(t) = x_0$ for all values of t .

Relation between the Pictures. The space of solutions and the space of initial data are not isomorphic — this is a general feature of theories with gauge symmetries. On the other hand, as we just noted, the reduced space of

¹⁵⁴More precisely: (q_1, π_1) and (q_2, π_2) belong to the same gauge orbit if and only if there is a solution g , instants $\Sigma_1, \Sigma_2 \subset (V, g)$, and diffeomorphisms $d_1 : S \rightarrow \Sigma_1$ and $d_2 : S \rightarrow \Sigma_2$ such that for $i = 1, 2$ (q_i, π_i) is the pull back to S by d_i of the initial data that g induces on Σ_i . Note that if $\Sigma_1 = \Sigma_2$ but $d_1 \neq d_2$, then (q_1, π_1) and (q_2, π_2) will be distinct but gauge-equivalent descriptions of the geometry of a single Cauchy surface in (V, g) .

¹⁵⁵See [Fischer and Moncrief, 1996].

¹⁵⁶If (q, π) is the geometry of a Cauchy surface in g , then canonical isomorphism between \mathcal{I}' and \mathcal{S}' sends $[q, \pi]$ to $[g]$.

solutions and the reduced space of initial data are believed to be isomorphic. In the case of a theory on a fixed background spacetime, a slicing yields a one-parameter family of symplectic isomorphism between the space of solutions and the space of initial data that serves the dual purposes of intertwining the temporal symmetries of their respective spaces and allowing us to construct a representation of change on the space of solutions. In the present case we have only a single canonical isomorphism between the two spaces.

Time. On neither the space of solutions nor reduced space of solutions do we find an action of the real numbers implementing time translation. Nor do we find a non-trivial action implementing time evolution on the reduced space of initial data, since the Hamiltonian trajectories are all trivial there. On the space of initial data, we do have non-trivial Hamiltonian trajectories. But a dynamical trajectory on the space of initial data cannot in general be viewed as encoding time evolution: there is nothing, for instance, to prevent such a trajectory from being periodic, even when the solution corresponding to the gauge orbit the trajectory lives in is not periodic in any sense.

There is, however, a class of dynamical trajectories on the space of initial data that can be viewed as encoding dynamics — those trajectories that correspond to sequences of initial data that could be stacked to form sensible spacetime geometries (when this is possible, the result is always a solution of the field equations). Through each point of the space of initial data there are in fact many such trajectories. But, as is usual in theories with gauge symmetry, there is no privileged way of cutting down this multitude to a distinguished subset that encode time evolution via an \mathbb{R} -action.

Change. Let us take some changeable physical quantity like the instantaneous spatial volume of the universe. How would we represent such a quantity on the various spaces in play? On both the space of solutions and the reduced space of solutions, we face our usual problem: points in these spaces represent history timelessly, so no function on such a space can represent in a direct way a changeable physical quantity. In the past, we were able to get around this problem using one of the following strategies. (i) We could find a function on f on a space arising on the Hamiltonian side, then use a slicing-dependent one-parameter family of isomorphisms between this space and the (reduced) space of solutions to find a one-parameter family of functions on the latter space encoding the behaviour of the given quantity. (ii) Or we could find a function on the (reduced) space of solutions encoding the value of the quantity of interest at a given instant, then use a dynamical time translation group on the (reduced) space of solutions to generate a one-parameter family of such functions. Neither of these strategies will work this time: we do not have a one-parameter family of isomorphisms indexed by instants, nor a notion of time translation on the (reduced) space of solutions.

We are in fact no better off on the reduced space of initial data: there too points correspond to entire histories of the system, and individual functions

are ill-suited to represent changeable quantities. And on the space of initial data we face an unattractive dilemma: if we seek to represent changeable quantities via non-gauge invariant functions, then we face indeterminism; if we employ gauge-invariant functions, then we are faced with essentially the same situation we met in the reduced space of initial data.

General Relativity in the Asymptotically Flat Regime

It is illuminating to consider a second sector of general relativity, in which one requires solutions to be asymptotically flat at spatial infinity. This case is of marginal physical interest, but it helps us to clarify the source of the problems we ran into in the spatially compact case.

In this regime our spacetime is \mathbb{R}^4 and kinematically possible fields are assignments of Lorentz signature metrics to V that are required to be, in an appropriate sense, asymptotically flat at spatial infinity.¹⁵⁷ Instants are also required to satisfy asymptotic conditions.

In this setting it is natural to consider $\mathcal{D}^\infty(V)$, the group of diffeomorphisms that leave the boundary conditions invariant, rather than the full group of diffeomorphisms. We find that the subgroup, $\mathcal{D}_0^\infty(V)$, of $\mathcal{D}^\infty(V)$ consisting of diffeomorphisms asymptotic to the identity at infinity is the largest group of gauge symmetries of the Lagrangian formulation of the theory and that $\mathcal{D}^\infty(V)$ is the semi-direct product of $\mathcal{D}_0^\infty(V)$ with the Poincaré group (every element of $\mathcal{D}^\infty(V)$ can be thought of as a product of an element of $\mathcal{D}_0^\infty(V)$ and a Poincaré symmetry acting at infinity).¹⁵⁸ The space of solutions of this theory carries a presymplectic form and breaks into gauge orbits, with two solutions in the same gauge orbit if and only if they differ by a diffeomorphism in $\mathcal{D}_0^\infty(V)$.¹⁵⁹ Diffeomorphisms in $\mathcal{D}_0^\infty(V)$ fix the gauge orbits; those in $\mathcal{D}^\infty(V)$ but not $\mathcal{D}_0^\infty(V)$ permute them. The significance of this is most clear at the level of the reduced space of solutions: this is a symplectic space carrying a representation of the Poincaré group — and in particular, for each notion of time translation at spatial infinity, this space carries a non-zero Hamiltonian generating this notion.¹⁶⁰

One can also give a Hamiltonian treatment of this sector of general relativity.¹⁶¹ One constructs the space of initial data as in the spatially compact case, except that conditions must be imposed on the asymptotic behaviour of the instantaneous

¹⁵⁷There are several notions of asymptotic flatness at spatial infinity. In this section, results are cited that are derived using three distinct but closely related approaches: (i) that of [Andersson, 1987]; (ii) that of [Ashtekar *et al.*, 1991]; and that of [Beig and Ó Murchadha, 1987]. For ease of exposition, I gloss over the differences in these approaches in the text — I do not believe that the result is misleading. For the relations between approaches (i) and (iii), see [Andersson, 1987, Definitions 2.3 and 2.4] and [Andersson, 1989, 78]. Both of approaches (ii) and (iii) are situated by their protagonists with respect to that of [Beig and Schmidt, 1982]; see [Ashtekar and Romano, 1992, §7] and [Beig and Ó Murchadha, 1987, §§4 and 5].

¹⁵⁸See [Andersson, 1987, Theorem 2.2] and [Ashtekar *et al.*, 1991, §3.3].

¹⁵⁹See [Ashtekar *et al.*, 1991, §3].

¹⁶⁰See [Andersson, 1987].

¹⁶¹See [Beig and Ó Murchadha, 1987].

field configuration and momentum. The resulting space carries a presymplectic form. Initial data sets (q, π) and (q', π') belong to the same gauge orbit if and only if there is a solution g and instants $\Sigma, \Sigma' \subset V$ such that Σ and Σ' are related by an element of $\mathcal{D}_0^\infty(V)$ and g induces (q, π) on Σ and (q', π') on Σ' .¹⁶² Just as on the space of solutions, we have a set of functions that can be viewed as the infinitesimal generators of the Poincaré group at infinity. Corresponding to a generator of time translations at infinity is a Hamiltonian on the space of initial data that generates a gauge equivalence class of notions of dynamics, each of which carries one from gauge orbit to gauge orbit (compare with the notion of dynamics on the space of initial data of an ordinary theory with gauge symmetries). So a generic dynamical trajectory, $x(t)$, generated by such a Hamiltonian will represent a nontrivial trajectory through the space of initial data; the same Hamiltonian will generate many trajectories through each point in the space of initial data; but each of these trajectories will agree for each value of t about the gauge orbit in which the state of the system dwells at that time.

One expects that the reduced space of initial data should be a symplectic space isomorphic to the reduced space of solutions and carrying a representation of the Poincaré group. Choosing a notion of time translation at infinity should pick out a Hamiltonian on the reduced space of initial data whose dynamical trajectories encode the dynamics of the theory: fixing a notion of the time translation, the corresponding Hamiltonian, and an arbitrary point in the reduced space of initial data, we should find that the Hamiltonian trajectory through this point encodes a sequence of equivalence class of instantaneous data, and that any way of picking representatives of these classes that stack to form a sensible spacetime geometry encodes a solution of the theory.

So the situation in this case is very different from that we saw above in the spatially compact case. We have representations of the Poincaré group on the reduced space of solutions and on the reduced space of initial data, and we have these representations encoded in structures on the space of solutions and the space of initial data.

And we can represent changeable physical quantities in a quite familiar way, via smooth functions on the reduced space of initial data. Special cases aside, such functions change their values as the state moves along the dynamical trajectories in that space. And we can calculate the rate of change of such quantities, etc. The situation is more complicated if we seek to represent change via functions on the space of solutions — this requires some of the apparatus to be developed below in section 7.3. But at least at the intuitive level, it is clear what needs to be done: because for each point in the reduced space of solutions, there is, for each notion of time translation at infinity, a one-parameter family of points in this space that correspond to the time translates of the given point, it ought to be possible to find, for any function on the reduced space of initial data that corresponds to a

¹⁶²As in the spatially compact case, distinct points in the space of initial data can correspond to the same field configuration and momentum induced by g on $\Sigma \subset V$, if we use distinct diffeomorphisms to pull back these tensors to the abstract instant S .

changeable quantity, a one-parameter family of functions on the reduced space of solutions that encode the value of that quantity at different moments of time.

Is General Covariance Special?

Einstein believed that the general covariance of general relativity was a very special feature with momentous physical consequences. Motivated by the observation that in special relativity there is a tight connection between the fact that the laws assume the same form in every inertial frame and the fact that all inertial observers are equivalent (so that there is no notion of absolute velocity), Einstein hoped that because the laws of his theory of gravity held in arbitrary coordinates the theory would be one in which *all* observers were equivalent (so that there would no notion of absolute motion whatsoever).

But, notoriously, the means were inadequate to the end: in general relativity there is a perfectly cogent (and coordinate-independent) distinction between those observers who are accelerated and those who are unaccelerated, between those who are rotating and those who are not.¹⁶³

Einstein's requirement that the laws of his theory should hold in arbitrary coordinate systems is just the translation into the language of coordinates of our first, weak, sense of general covariance. The preceding paragraph points out that this requirement does not have the powerful consequences that Einstein believed it to. Even worse, it was pointed out already by Kretschmann in 1917 that this weak sense of general covariance is not a very unusual feature: many pre-general relativistic theories can be given a weakly generally covariant formulation.¹⁶⁴ Indeed, there is a recipe that takes as input a Lagrangian field theory on a fixed background spacetime and gives as output a strongly generally covariant reformulation/relative of the given theory.¹⁶⁵

EXAMPLE 43 (Artificial General Covariance). Let T_0 be the theory of a massless Klein–Gordon scalar field, Ψ propagating on a fixed background spacetime, (V_0, g_0) . The Lagrangian for T_0 is $L_0(\Psi) := \frac{1}{2}g_0^{ab}\nabla_a\Psi\nabla_b\Psi$ and the corresponding equation of motion is $\square_0\Psi = 0$, where \square_0 is the d'Alembertian corresponding to g_0 .¹⁶⁶ Given T_0 we can construct a strongly generally covariant theory T as follows.¹⁶⁷ Let V be a manifold diffeomorphic to V_0 . The spacetime of T is the bare

¹⁶³Einstein's line of thought founders on the following observation: in special relativity Lorentz transformations are symmetries of the spacetime metric that is used to determine the state of motion of an observer, in general relativity an arbitrary diffeomorphism is certainly not a symmetry of the spacetime geometry of a given solution — but this geometry again plays a role in determining the state of motion of an observer. See [Friedman, 1983, Chapters II and V].

¹⁶⁴On Kretschmann, see [Rynasiewicz, 1999].

¹⁶⁵It is not obvious how one should individuate theories in the present context. For discussion and suggestions, see [Sorkin, 2002, 698] and [Earman, 2006, §4].

¹⁶⁶The d'Alembertian corresponding to a Lorentz metric g is defined just as the Laplacian of a Riemannian metric g : as $\text{div}_g \circ \text{grad}_g$ where div_g is the divergence operator of g and grad_g is the gradient operator of g .

¹⁶⁷See [Lee and Wald, 1990, 734] or [Torre, 1992, §II]. The same procedure will work for any scalar field with a first-order Lagrangian that features a non-derivative coupling of the field to

manifold, V , unequipped with any geometry. T involves two fields, X and $\Phi : X$ takes values in V_0 while Φ takes values in \mathbb{R} . A pair (X, Φ) counts as kinematically possible only if $X : V \rightarrow V_0$ is a diffeomorphism.¹⁶⁸ The Lagrangian L of T is constructed as follows: for any kinematically possible (X, Φ) , the n -form $L(X, \Phi)$ on V is the pullback to V by X of the n -form $L_0(\Psi)$ on V_0 , where $\Psi := \Phi \circ X^{-1}$. L admits $\mathcal{D}(V)$ as a group of gauge symmetries — so T is strongly generally covariant. Note that a kinematically possible pair (X, Φ) is a solution of T if and only if $\Psi = \Phi \circ X^{-1}$ is a solution of T_0 . This is equivalent to saying that a pair (X, Φ) is a solution if and only if Φ is a solution of the massless Klein–Gordon equation $\square\Phi = 0$, with \square the d’Alembertian corresponding to the metric $g := X^*g_0$ on V .

This shows that there are relatively ordinary theories, like the theory of Klein–Gordon field, that can be given strongly generally covariant formulations. So even strong general covariance fails to distinguish general relativity from perfectly pedestrian theories.

Nonetheless, it is difficult to shake the feeling that the special nature of general relativity among physical theories has something to do with its general covariance. Indeed, it would appear that at the present time the best that can be said is that what makes general relativity special is that its most natural and perspicuous formulations are generally covariant. But that is just to say that we do not yet understand the matter, I think.

In this connection, it is natural to ask whether the difficulties that we encounter in representing time and change in general relativity arise for the artificially strongly generally covariant theory of example 43.

EXAMPLE 44 (Artificial General Covariance and the Problem of Time). Let us return to the theories T_0 and T of example 43, and let us assume for convenience that the spacetime, (V_0, g_0) , of T_0 does not admit any isometries. Suppose that we were simply handed T . Would there be any way to represent changeable quantities via functions on a symplectic space associated with T ?

Let \mathcal{S} be the space of solutions of T , and let \mathcal{S}' be the corresponding reduced space (i.e., the space of gauge orbits of \mathcal{S}). As one would expect, two solutions (X, Φ) and (X', Φ') lie in the same gauge orbit of \mathcal{S} if and only if there exists $d \in \mathcal{D}(V)$ such that $X' = X \circ d$ and $\Phi' = \Phi \circ d$. The space \mathcal{S} is, of course, presymplectic while the space \mathcal{S}' is symplectic. But since solutions in the same gauge orbit will not agree about the value of Φ or X at any point of V , it is difficult to view a diffeomorphism equivalence class of solutions as assigning properties to points of V , and so it would appear to be impossible to think of \mathcal{S}' as the space of solutions corresponding to some local Lagrangian. By following a procedure like that used in the discussion of general relativity above, we can construct the space of initial data, \mathcal{I} , of T , and the corresponding reduced space, \mathcal{I}' . The latter will be a symplectic space. But note that the Hamiltonians on \mathcal{I} and \mathcal{I}' vanish. So

the spacetime metric.

¹⁶⁸Strictly speaking, this takes us outside of our official framework for Lagrangian field theories, since the value that X takes at distinct points of V are not independent of one another.

although we have been able to construct symplectic spaces, we do not have the nontrivial flows associated with time translation or time evolution that we require to set up our representation of change via functions on these spaces. So far, the present case looks very much like the case of spatially compact general relativity.

But now note that from knowledge of T alone we can reconstruct T_0 . The field X has as its target space the manifold V_0 . We take T_0 to be the theory of a scalar field Ψ on V_0 with Lagrangian L_0 given as follows: let Ψ be a kinematically possible field of T_0 and let $X : V \rightarrow V_0$ be an arbitrary diffeomorphism; then we define $L_0(\Psi)$ to be the n -form on V_0 that results when we use X^{-1} to pullback to V_0 the n -form $L(X, \Psi \circ X)$; the result is independent of the X chosen. The resulting equations of motion is $\square_0 \Psi = 0$. Noting that \square_0 arises as the d'Alembertian of a unique metric g_0 on V_0 and that field propagates causally relative to g_0 , it is natural for us to view g_0 as the geometrical structure of V_0 , and go on to consider slicings relative to g_0 , etc.

With T_0 in hand, we can construct the space of solutions, \mathcal{S}_0 . Relative to a slicing of (V_0, g_0) , we can represent any changeable quantity — e.g., the volume of the support of the scalar field — via functions on \mathcal{S}_0 in the usual way.

Finally, note that \mathcal{S}' is canonically symplectically isomorphic to \mathcal{S}_0 .¹⁶⁹ So we can transfer our representation of change from the latter space to the former. So there is a way to avoid the problem of time in this case.¹⁷⁰

There is, however, an obvious worry about this approach. Let \tilde{g}_0 be a metric on V distinct from g_0 . Then \square_0 is not the d'Alembertian of \tilde{g}_0 ; but presumably this operator is still definable in terms of \tilde{g}_0 . So according to \tilde{g}_0 the Euler–Lagrange equations of L_0 on V_0 are not the Klein–Gordon equations, but some less famous equations. Now, the above strategy amounts to thinking of T as *really* the theory of a Klein–Gordon field on a spacetime isomorphic to (V_0, g_0) . But it was no part of our data that T is a Klein–Gordon theory. So what is to stop us from thinking of T as *really* a theory of a field obeying some less famous equations on a spacetime isomorphic to (V_0, \tilde{g}_0) ? In this case we would use slicings of (V_0, \tilde{g}_0) to set up our representation of change, etc.

Here are two things one might say in response to this worry. (1) We sought and found a natural way of representing change via functions on \mathcal{S}' . It is no problem if there are others. (2) We normally demand that of a physically reasonable theory that its field propagate along the nullcones of the spacetime metric. This will be true of T_0 only for metrics \tilde{g}_0 conformally equivalent to g_0 .¹⁷¹ Every slicing of (V_0, g_0) is also a slicing of (V_0, \tilde{g}_0) for each \tilde{g}_0 conformally related to g_0 (since

¹⁶⁹Via the map that sends an equivalence class, $[X, \Phi]$, of solutions T to the solution $\Psi = \Phi \circ X^{-1}$ of T_0 . It is at this point that we require the assumption the g_0 does not admit isometries: in general, \mathcal{S}' is isomorphic to the quotient of \mathcal{S}_0 by the action of the isometry group of g_0 .

¹⁷⁰Note that we must choose a slicing of (V_0, g_0) in order to get a family of functions on \mathcal{S}' corresponding to a changeable physical quantity. Such functions tell us things like how large the volume of the support of the field is at the instant when the geometry of space assumes a given form.

¹⁷¹Recall that metrics g_0 and g_1 on V are conformally related if there is a positive scalar $\Omega : V \rightarrow \mathbb{R}$ such that $g_1 = \Omega g_0$.

conformally related methods agree about which lines are timelike and which hypersurfaces are spacelike). So relative to such a slicing we can consider a quantity that is conformally invariant in the sense that for each $\Sigma \subset V$ in our slicing, this quantity has the same value on Σ in (V_0, \tilde{g}_0, Ψ) for each \tilde{g}_0 conformally related to g_0 . Such a quantity is represented by the same one-parameter family of functions on the reduced space of solutions of T whether we view T as secretly a theory of a Klein–Gordon field on a spacetime isomorphic to (V_0, g_0) or as secretly a theory of some other sort of field on a spacetime isomorphic to (V_0, \tilde{g}_0) .

7.2 The Problem of Time

In each of the theories considered in sections 5 and 6 above, the dynamical content of the theory was encoded in a flow (possibly time-dependent, possibly merely local) on a symplectic space of states within the Lagrangian or Hamiltonian formulation of the theory. That this fails in general relativity, conceived of as a theory of the universe as a whole, is what sets that theory apart. And, of course, this feature means that the standard strategies for representing change also fail for this theory: since one does not have a flow corresponding to time evolution on the reduced space of initial data, no function on that space can represent a changeable physical quantity; it follows that one does not have the apparatus required to represent changeable quantities via functions on the reduced space of solutions either.

This nexus is the problem of time: time is not represented in general relativity by a flow on a symplectic space and change is not represented by functions on a space of instantaneous or global states.¹⁷²

Before proceeding to discuss the significance of this problem it is important to be clear about its nature and sources.

- If one approaches the problem of time via a focus on the transition from the space of initial data to the reduced space of initial data, the problem can appear especially urgent. For in passing from the space of initial data to the reduced space of initial data, one identifies initial data sets that correspond to distinct Cauchy surfaces within a single solution. *Prima facie*, this involves treating the current state of the universe and its state just after the Big Bang as the *same state*. Moral: according to general relativity, change is an illusion.

But this is too hasty. For of course the reduced space of initial data is canonically isomorphic to the reduced space of solutions.¹⁷³ And in this latter space, some points represent worlds in which there is change (e.g., worlds which begin with a Big Bang) and some represent changeless worlds

¹⁷²The canonical presentations of the problem of time are [Kuchař, 1992] and [Isham, 1993]. For philosophical discussions, see [Belot and Earman, 2001], [Butterfield and Isham, 2000], and [Earman, 2002]. For critical reactions to this literature, see [Maudlin, 2002] and [Healey, 2002].

¹⁷³Under the map that sends $[q, \pi]$ to $[g]$ if (q, g) describes the instantaneous state on some Cauchy surface of (V, g) .

(e.g., world modelled by Einstein's static solution). So it is hard to see how general relativity teaches us the moral announced.

So I would like to disavow formulations of the problem of time that rely on this way of speaking. More constructively, I would like to suggest that it is helpful to concentrate on the reduced space of solutions rather than on the reduced space of initial data in setting up the problem of time. In the well-behaved theories of section 5 the space of initial data and the space of solutions are symplectically isomorphic, but we nonetheless think of these two spaces as having distinct representational functions — roughly and heuristically speaking, one is suited to represent possible instantaneous states while the other is suited to represent possible worlds. This distinction is grounded by the fact that relative to a slicing one finds that for each $t \in \mathbb{R}$, the map T_{Σ_t} that sends a solution to the initial data that it induces on the instant $\Sigma_t \subset V$ defines a distinct isomorphism between the space of solutions and the space of initial data. This makes it natural to think of points of the latter space as representing states (universals) that can occur at distinct times and to think of points in the space of solutions as representing possible worlds composed out of such states. The elements of this story survived more or less unscathed the introduction of various complicating factors in section 6. But in the case of cosmological general relativity we have only a single canonical isomorphism between the reduced space of initial data and the reduced space of solutions. In this context, it is difficult to deny that the reduced space of solutions and the reduced space of initial data are representationally equivalent. And it seems straightforward that we should interpret points in the reduced space of solutions as representing general relativistic worlds rather than instantaneous states — so we should say that same thing about points in the reduced space of initial data. Thus, we should resist any temptation to think of the reduction procedure as telling us to think of an early state of the universe and a late state of the universe as being the same instantaneous state.

- Since we have been focussing on the Lagrangian rather than the Hamiltonian picture, but have nonetheless run straight into the problem of time, we can conclude that this problem is not an artifact of the 3+1 decomposition involved in the Hamiltonian approach. Likewise, the problem of time is a feature of general relativity as a cosmological theory, but not of general relativity in the regime of asymptotic flatness at spatial infinity, nor of field theories on fixed relativistic backgrounds, nor, I think, of the artificial strongly generally covariant theory of examples 43 and 44 above. From this we can conclude that the following are *not* sufficient conditions for the problem of time: the lack of a preferred slicing; the jiggleability of admissible slicings; the invariance of the theory under a group of spacetime diffeomorphisms. It appears that the problem arises when we employ a diffeomorphism-invariant theory to model a situation in which we take geometry to be fully dynamical

(i.e., we do not smuggle in any background structure, at spatial infinity or elsewhere).

For everything that I have said so far, the Problem of Time may sound like no more than a diverting puzzle. Granted, time does not appear as a symmetry in general relativity as it did in earlier theories (even in the infinitesimal sense involved in a local flow). But, of course, part of the allure of the theory is that it changes the nature of time in a fundamental way. And since successful applications of the theory involve the representation of changeable physical quantities (e.g., the perihelion of Mercury), it would seem that there *must* be some way of way of generalizing the picture of the previous sections to cover general relativity. And while it will be granted that a search for this generalization might turn out to be enlightening, it may well not seem a very pressing project.

This puzzle begins to look far more urgent when we turn our attention to quantization. The good news is that upon reduction, one ends up with a symplectic space representing the true degrees of freedom of general relativity. Without something along these lines, quantization would be impossible. But the vanishing of the Hamiltonian for cosmological general relativity means that two looming difficulties block the road to the successful quantization of general relativity.

1. What is one to do next? Normally a Hamiltonian or a Lagrangian plays a crucial role in quantization. One defines quantum dynamics via these objects. In the case of spatially compact general relativity the reduced space of initial data inherits from the original space of initial data a Hamiltonian — which vanishes, so that the corresponding dynamics is trivial. And it does not appear to make any sense to speak of a local Lagrangian field theory of the true degrees of freedom of the gravitational field. The way forward is unclear.
2. Furthermore, it is not clear how one would make sense of a quantization of general relativity. While in the classical theory one can find change in solutions even without being able to find it at the dynamical level (in terms of quantities on the space of solutions, etc.), it is not obvious how this could be done at the quantum level. Perhaps the best that one can hope for is to be able to speak of approximate time and change in a subset of quantum states that approximate classical solutions. That seems perfectly acceptable — what one should be aiming at, even, in a theory in which the geometry of space and time are themselves quantized. But the usual techniques of semi-classical approximation require a Hamiltonian.¹⁷⁴

¹⁷⁴Thus, the WKB method aims to construct approximate eigenstates for the quantum Hamiltonian. Analyses based upon decoherence, coherent states, etc., aim to show that the dynamics driven by the quantum Hamiltonian approximates that of the corresponding classical system, and so on. See [Landsman, this volume].

7.3 Finding Time in General Relativity

This final section discusses what is probably the most obvious way around the problem of time. In the cases discussed in sections 5 and 6, we were able to represent change via functions on the (reduced) space of solutions of the theory because we had a slicing, $\sigma : S \times \mathbb{R} \rightarrow V$, that decomposed spacetime into space and time, and thereby allowed us to identify functions on the (reduced) space of solutions that corresponded to the values of a given quantity at different instants. But the notion of a solution-independent decomposition of spacetime into space and time makes no sense in general relativity, since solutions differ as to which curves count as timelike and which hypersurfaces count as spacelike. Somewhat surprisingly, it turns out to be possible to construct a Hamiltonian version of general relativity without employing slicings. But — unsurprisingly — without some sort of decomposition of spacetime into instants, it makes no sense to ask which states follow a given state (so there is no real dynamics on the Hamiltonian side) nor to try to construct a one-parameter family of functions on the reduced space of solutions that corresponds to the instantaneous values of a quantity of interest. So it is natural to look for a surrogate of the notion of a slicing that applies to diffeomorphism equivalence classes of solutions, rather than to individual solutions — and to hope that this will lead to familiar-looking accounts of the representation of time and change.

Throughout this final subsection, unless otherwise noted, I restrict attention to spatially compact vacuum general relativity in four spacetime dimensions with vanishing cosmological constant.

Let me begin with some definitions.

DEFINITION 45 (Geometry). A point in the reduced space of solutions of general relativity is called a *geometry*. A geometry is an orbit of the action on the space of solutions of the group $\mathcal{D}(V)$ of diffeomorphisms from V to itself. We write $[g]$ for the geometry corresponding to a solution g ; we speak of a solution in $[g]$ as a solution *with geometry* $[g]$.

DEFINITION 46 (Instantaneous Geometry). The group $\mathcal{D}(S)$ of diffeomorphisms from the abstract instant S to itself acts on the space of initial data. We call an orbit of this action an *instantaneous geometry*. We denote the instantaneous geometry corresponding to an initial data point (q, π) by $\langle q, \pi \rangle$. We speak of initial data (q, π) as having the geometry $\langle q, \pi \rangle$.¹⁷⁵

DEFINITION 47 (Time for a Solution). Let (V, g) be a solution. A *time for* (V, g) is a partition, $\{\Sigma\}$, of (V, g) by Cauchy surfaces, called the *instants* of the time. A *parameterized time* is a time together with a preferred parameterization of the set of instants. An *affinely parameterized time* is a time whose instants are

¹⁷⁵Note that an instantaneous geometry is *not* a point in the reduced space of initial data: initial data induced by a given solution on distinct Cauchy surfaces correspond to the same point in the reduced space of initial data, but (in general) to distinct points in the space of instantaneous geometries.

parameterized up to the choice of origin.¹⁷⁶

DEFINITION 48 (Absolute Time). Let g be a solution. A time $\{\Sigma\}$ for (V, g) is called *absolute* if every isometry of g maps instants in $\{\Sigma\}$ to instants in $\{\Sigma\}$. An affinely parameterized time for g is called *absolute* if each isometry of g satisfies the preceding condition and preserves the parameter difference between each pair of instants. A parameterized time for g is called *absolute* if each isometry of g maps each instant of the time to itself.

Every globally hyperbolic solution admits a parameterized time (since each globally hyperbolic solution can be foliated by Cauchy surfaces, which can be given an arbitrary parameterization). But it is not always possible to find absolute times for solutions with large isometry groups. Minkowski spacetime does not admit an absolute time.¹⁷⁷ If a spacetime admits time translation or inversion as a symmetry, then it does not admit an absolute parameterized time.

DEFINITION 49 (Time for General Relativity). A (*plain, affinely parameterized, or parameterized*) *time for general relativity* is a map defined on a subset of the space of solutions that assigns to each solution in its domain a (plain, affine, or parameterized) time for that solution, and does so in an appropriately smooth manner.

DEFINITION 50 (Geometric Time for General Relativity). A (plain, affinely parameterized, or parameterized) time for general relativity is called *geometric* if it satisfies the following conditions. (i) Its domain of definition is closed under the action of $\mathcal{D}(V)$ on the space of solutions. (ii) If g and g' are in the domain of the time and $g' = d^*g$ for some diffeomorphism $d : V \rightarrow V$, then the foliation assigned to g' is the image under d^{-1} of the foliation assigned to g (if the time is affinely parameterized, then we require that such a d preserve the time difference between any two instants; if the time is parameterized, then we require that such a d map the instant labelled by t to the instant labelled by t). I will often shorten *geometric time for general relativity* to *geometric time*.

REMARK 51 (Geometric Times are Absolute). The (plain, affinely parameterized, or parameterized) time that a geometric time for general relativity assigns to a solution g is always absolute. For if $d : V \rightarrow V$ is an isometry of g , then condition (ii) in the preceding definition tells us that d preserves the time assigned to g , together with its parameterization properties, if any. It follows that Minkowski spacetime is not in the domain of definition of any geometric time for general relativity, and that no solution invariant under time translation or inversion is in the domain of definition of any parameterized geometric time for general relativity.

We can think of a (parameterized, affinely parameterized, or unparameterized)

¹⁷⁶We can think of a time for (V, g) as an unparameterized curve in the space of Cauchy surfaces of (V, g) ; a parameterized time is a parameterized curve of this type; an affinely parameterized time is an affinely parameterized curve of this type.

¹⁷⁷A time invariant under the notion of time translation associated with a given frame will fail to be invariant under boosts relative to that frame. The same argument will work in de Sitter spacetime, or in other spacetimes admitting boost symmetries; see [Moncrief, 1992] for examples.

geometric time for general relativity as a means of associating a geometry $[g]$ in the reduced space of solutions with a (parameterized, affinely parameterized, or unparameterized) curve $\langle q(t), \pi(t) \rangle$ in the space of instantaneous geometries; we call such a curve a *dynamical trajectory*. The correspondence between geometries and dynamical trajectories is set up in the obvious way: let g be a solution in the domain of definition of a given geometric time, and let $(q(t), \pi(t))$ be the (parameterized, affinely parameterized, or unparameterized) curve in the space of initial data that results when we look at the initial data induced by g on the instants in the time assigned to g ; $\langle q(t), \pi(t) \rangle$ is the dynamical trajectory we seek.¹⁷⁸ If g_1 and g_2 are solutions with the same geometry, then they are related by some diffeomorphism $d : V \rightarrow V$. In this case d also relates the foliations assigned to them by our geometric time, so g_1 and g_2 will correspond to the same dynamical trajectory in the space of instantaneous geometries.

A number of interesting examples of geometric times are known. Most have very small domains of definition: (i) within the class of nonrotating dust solutions, a geometric time is given by foliating each solution by the unique family of hypersurfaces everywhere orthogonal to the dust worldlines; (ii) within the class of solutions whose isometry groups are three-dimensional with spacelike orbits, a geometric time is given by foliating each solution by the orbits of its isometry group.¹⁷⁹ Examples of wider scope are harder to come by but do exist.

EXAMPLE 52 (CMC Time). Recall that if $\Sigma \subset V$ is a Cauchy surface for (V, g) , then we can define tensors q^{ab} and k_{ab} on Σ with the following meaning: $q^{ab} := g^{ab} \lfloor_{\Sigma}$ is the Riemannian metric that g induces on Σ and $2k_{ab}$ is the rate of change of this metric according to freely falling observers whose worldlines intersect Σ orthogonally. Out of these tensors we can construct the *mean curvature*, $k : \Sigma \rightarrow \mathbb{R}$, defined by $k := q^{ab}k_{ab}$ (so $k(x)$ is just the trace of the matrix that encodes information about k_{ab} at x). A Cauchy surface $\Sigma \subset V$ for a solution g is called a *surface of constant mean curvature*, or simply a *CMC surface*, if k is a constant function on Σ . Recall that unless otherwise noted, we restrict attention to (3+1) spatially compact globally hyperbolic vacuum solutions with vanishing cosmological constant.

Applicability. It is widely believed that a large class of solutions to Einstein's field equations can be foliated by CMC surfaces.

1. It is known that the set of solutions containing a CMC slice is an open set in the space of solutions.¹⁸⁰

¹⁷⁸Strictly speaking, in order to construct the curve $(q(t), \pi(t))$ in the space of initial data, we need to introduce a slicing of (V, g) whose instants coincide with those of the given time, so that we can pullback states on concrete instants to states on our abstract instant S ; the arbitrariness involved in a choice of slicing washes out when we quotient the space of initial data by the action of $\mathcal{D}(S)$ to reach the space of instantaneous states.

¹⁷⁹Scheme (i) generalizes Einstein's simultaneity convention to the context of dust cosmology; see [Sachs and Wu, 1977, §5.3]. Note that schemes (i) and (ii) need not coincide within their shared domain of definition; see [King and Ellis, 1973].

¹⁸⁰See, e.g., [Isenberg and Marsden, 1982, 195].

2. It was once conjectured that all solutions contain at least one CMC surface, but it is now known that this is not so.¹⁸¹
3. It is was once conjectured that all solutions admitting a CMC slice can be foliated by such slices.¹⁸² This is now believed to hold only for certain spatial topologies.¹⁸³
4. It is believed that within the class of solutions foliated by CMC slices, all solutions of a given spatial topology will exhibit the same range of values of constant mean curvature, with the only exceptions being stationary solutions (recall that a solution is stationary if it admits a timelike Killing field — roughly speaking, the infinitesimal generator of a time translation group).¹⁸⁴

Invariance Properties. CMC foliations behave superbly well under isometries.¹⁸⁵

Let (V, g) be a solution, $\{\Sigma\}$ a set of CMC surfaces that foliates V , and $d : V \rightarrow V$ an isometry of g . Then d leaves the foliation $\{\Sigma\}$ invariant.¹⁸⁶ If (V, g) is non-stationary, then: (a) any symmetry d of g preserves each leaf in $\{\Sigma\}$; and (b) for any real number κ , there is at most one Cauchy surface with constant mean curvature κ . If (V, g) is stationary then: g is flat and any CMC surface in (V, g) has vanishing mean curvature.¹⁸⁷

CMC Time. Foliating each solution by its CMC slices, when possible, determines a geometric time within the class of solutions we are considering. We can render this an affinely parameterized geometric time as follows: for non-stationary solutions, the parameter difference between slices of mean

¹⁸¹See [Bartnik and Isenberg, 2004, 32] or [Chruściel *et al.*, 2005, Corollary 1.3]. The corresponding conjecture for spatially compact dust solutions is also false; see [Bartnik, 1988].

¹⁸²For the original form of the conjecture, see, e.g., [Isenberg and Marsden, 1982, Conjecture 3.2]. This conjecture is known to be true for flat spacetimes ([Barbot, 2005, §12]) and the corresponding conjecture is known to be true in the (2+1) case ([Andersson *et al.*, 1997]). The counterpart of this conjecture is known to be false for spatially compact dust solutions ([Isenberg and Rendall, 1998]) and in the asymptotically flat vacuum case, where the Schwarzschild solution provides a counterexample ([Eardley and Smarr, 1979, §III]).

¹⁸³For the current conjecture, see [Rendall, 1996, Conjectures 1 and 2]. It is now believed that for some spatial topologies, behaviour analogous to that of the Schwarzschild solution can occur; see [Rendall, 1996] and [Andersson, 2004, 81]. In the (3+1) case, the revised conjecture is known to be true for some types of highly symmetric solutions, even when some forms of matter are allowed; see [Rendall, 1996, Theorems 1 and 2], [Andersson, 2004, 81 f. and 95], and the references therein.

¹⁸⁴For this, see [Rendall, 1996, Conjectures 1 and 2]. For the situation in highly symmetric cases and in (2+1) dimensions, see the references of the previous two footnotes.

¹⁸⁵See [Isenberg and Marsden, 1982, §3].

¹⁸⁶This would fail for spacetimes admitting boost symmetries, such as Minkowski spacetime and de Sitter spacetime. (Note that since we require vanishing cosmological constant, de Sitter spacetime does not count as a spatially compact vacuum solution for present purposes.) Note that in the asymptotically flat case, the question of the invariance of CMC slices is much more involved; see [Bartnik *et al.*, 1990, §5].

¹⁸⁷Of course, in general a timelike Killing vector does not guarantee flatness. But it does so within the class solutions presently under consideration.

curvature κ_1 and κ_2 is $|\kappa_2 - \kappa_1|$; for stationary solutions, the parameter difference between two slices is the proper time elapsed between those slices. If we restrict attention to non-stationary solutions, and assign to each slice the parameter value given by its mean curvature, then we arrive at a parameterized geometric time.

EXAMPLE 53 (Cosmological Time). Given a solution (V, g) , the *cosmological time function* for g is the map $\tau : V \rightarrow \mathbf{R} \cup \{\infty\}$ that assigns to each $x \in V$ the supremum over the length of all past-directed causal curves starting at x . Obviously there are many well-behaved spacetimes in which $\tau(x)$ is badly behaved — e.g., in Minkowski spacetime, $\tau(x) = \infty$ for all events. We say that the cosmological time function of a solution is *regular* if: (a) $\tau(x) < \infty$ for all x and (b) $\tau \rightarrow 0$ along each past inextendible causal curve. If τ is a regular cosmological time function on (V, g) then: (i) (V, g) is globally hyperbolic; (ii) τ is a time function for the solution in the usual sense (i.e., it is continuous and strictly increasing along future-directed causal curves); and (iii) the level surfaces of τ are future Cauchy surfaces in (V, g) (i.e., these surfaces have empty *future* Cauchy horizons).¹⁸⁸ In spatially compact vacuum (2+1)-dimensional general relativity, it is known: (a) that the cosmological time is regular for almost all spacetime topologies; and (b) that in one important class of solutions the cosmological time coincides with the CMC time.¹⁸⁹ On the class of spacetimes with regular cosmological time functions whose level surfaces are Cauchy surfaces, we construct a geometric time for general relativity by foliating each solution by the surfaces of constant cosmological time; parameterizing these foliations by the value the cosmological time function takes on each leaf yields a parameterized geometric time for general relativity, so long as we exclude solutions with a time reflection symmetry.

A geometric time for general relativity is, in effect, a means of separating out from the infinite number of variables of the theory one relative to which the others are to be seen as evolving, by allowing us to pass from a point in the reduced space of solutions to a (possibly unparameterized) curve in the space instantaneous geometries. Instantaneous physical quantities such as the volume of the universe, or the number of stars, or the size of the solar system can be represented by real-valued functions on the space of instantaneous geometries. And so the choice of a geometric time allows us to talk about change in the familiar way: we can check to see whether a function on the space of instantaneous geometries that represents a quantity of interest takes on different values at points corresponding to the different instantaneous geometries that occur in a given spacetime geometry. If our geometric time for general relativity is affinely parameterized, we can calculate the rate of change of quantities of interest (since we then have an affinely parameterized curve through the space of instantaneous geometries corresponding to a

¹⁸⁸See [Andersson *et al.*, 1998, Propositions 2.2 and 2.5 and Corollary 2.6].

¹⁸⁹See [Benedetti and Guadagnini, 2001]. In general, however, surfaces of constant cosmological time are less smooth than CMC surfaces, so the two notions of time do not coincide; see [Benedetti and Guadagnini, 2001, 331] or [Barbot and Zeghib, 2004, §5.4.1].

given spacetime geometry). If we have a parameterized geometric time for general relativity, we can even mimic the construction we used in earlier sections to represent changeable quantities by one-parameter families of functions on the reduced space of solutions of the theory: given the function f on the space of instantaneous geometries that represents our quantity of interest, and a real number t , we define a partially defined function f_t on the reduced space of solutions by setting $f_t[g]$ equal to the value that f takes on the instantaneous geometry corresponding to t in $[g]$.

As delineated above, the problem of time in general relativity had two major aspects.

1. Time is not represented in spatially compact general relativity, as it was in earlier theories, via a flow on a symplectic space of states;
2. Change is not represented, as it was in earlier theories, via functions on symplectic spaces corresponding to the spaces of possible instantaneous states and worlds.

We now see that if we go as far as introducing a parameterized geometric time, we can address the second of these worries by representing a changeable quantity by a one-parameter family of functions on the reduced space of solutions, in the usual way.

Does the introduction of a geometric time suffice to address the first worry? Any geometric time singles out a subspace of the space of instantaneous geometries, consisting of those $\langle g, \pi \rangle$ that arise as instantaneous geometries of the Cauchy surfaces picked out by that geometric time — e.g., only instantaneous geometries portraying space as having constant mean curvature can arise according to the CMC slicing scheme. If we introduce an affinely parameterized geometric time for general relativity, then we do get a flow on the space of instantaneous geometries that arise according to this geometric time (since this space is partitioned by the affinely parameterized dynamical trajectories corresponding to geometries in the domain of definition of the given geometric time). But one does not expect this space to be symplectic nor to be isomorphic to the reduced space of solutions.¹⁹⁰ So a flow on the space of instantaneous geometries associated with our geometric time is not a flow on a symplectic space. And since each dynamical trajectory on the space of instantaneous geometries corresponds to a single point in the reduced space of solutions, we have no means of carrying our flow on the former space over to a nontrivial flow on the latter.

A natural strategy to set up a representation of time via a flow on a symplectic manifold is to attempt to parlay a choice of geometric time for general relativity

¹⁹⁰Intuitively, the space of instantaneous geometries that arise according to a given geometric time can be thought of as the product of the reduced space of solutions with the real line (since each geometry corresponds to a one-parameter family of instantaneous geometries relative to the geometric time). So the space of instantaneous geometries of the given geometric time is not isomorphic to the space of solutions — nor can it be symplectic, since it is the product of a symplectic space with an odd-dimensional space.

into a reformulation of of the theory as a nontrivial (but possibly time-dependent) Hamiltonian system. In one important case, it is known that this can be achieved.

EXAMPLE 54 (CMC dynamics). We consider the CMC time introduced in example 52 above.¹⁹¹ We impose restrictions on the topology of our abstract instant S .¹⁹² Let \mathcal{M} be the space of Riemannian metrics on S with constant scalar curvature -1 .¹⁹³ The cotangent bundle $T^*\mathcal{M}$ is a symplectic space; an element of $T^*\mathcal{M}$ is of the form (γ, p) where $\gamma \in \mathcal{M}$ and p is a symmetric contravariant tensor density of rank two on S that is divergenceless and traceless according to γ . We consider $(\gamma, p), (\gamma', p') \in T^*\mathcal{M}$ to be equivalent if there is a diffeomorphism $d : S \rightarrow S$ such that $(\gamma', p') = (d^*\gamma, d^*p)$. The space $\mathcal{I}^* := T^*\mathcal{M}/\mathcal{D}(S)$ that results when we quotient out by this equivalence relation inherits a symplectic structure from $T^*\mathcal{M}$. We will call points in \mathcal{I}^* *conformal initial data*. For each $t < 0$ there is a geometrically natural symplectic isomorphism between \mathcal{I}^* and the space of instantaneous geometries with constant mean curvature t .¹⁹⁴ And there is a natural symplectic isomorphism between the latter set and the reduced space of solutions of general relativity (under which an instantaneous state is sent to the unique geometry that it occurs in). So for each $t < 0$ we have a symplectic isomorphism between the space of conformal initial data and the reduced space of solutions.

Conversely, given a geometry $[g]$ and a time $t < 0$ we can look for the point in \mathcal{I}^* that corresponds to $[g]$ according to the isomorphism labelled by t . Doing this for each $t < 0$ gives us a curve in \mathcal{I}^* corresponding to $[g]$. A generic point in \mathcal{I}^* will lie on many such trajectories: in general if $x \in \mathcal{I}^*$ and $t_1 \neq t_2$ then the instantaneous geometry of constant mean curvature t_1 corresponding to (x, t_1) and the instantaneous geometry of constant mean curvature t_2 corresponding to (x, t_2) will reside in different spacetime geometries. If we look at the complete family of trajectories in \mathcal{I}^* corresponding to all of the geometries in the reduced space of solutions, then we find that these are generated by the symplectic structure of \mathcal{I}^* together with a time-dependent Hamiltonian $h(t)$ that is a simple function of t

¹⁹¹For an overview of the (3+1) case, see [Fischer and Moncrief, Unpublished, §§2 and 3]; for details see [Fischer and Moncrief, 1996], [Fischer and Moncrief, 1997], and the references therein. For the (2+1) case see [Moncrief, 1989] and [Andersson *et al.*, 1997]. The construction described below is an example of deparameterization. For this notion and for finite-dimensional applications, see [Beig, 1994, §2].

¹⁹²We impose two conditions. (i) S must be of Yamabe type -1 , i.e., the only constant scalar curvature Riemannian metrics that S admits have negative scalar curvature. This is essential for the constructions employed in the papers cited. (ii) S must not admit any Riemannian metrics with isometry groups of positive dimension. This saves us from having to worry about singular quotient spaces.

¹⁹³Because S is of Yamabe type -1 , every Riemannian metric on S is conformally equivalent to a metric in \mathcal{M} .

¹⁹⁴Let us ignore the $\mathcal{D}(S)$ symmetry for a moment. Given a pair (γ, p) and a time $t < 0$ there is a unique positive scalar ϕ on S solving the Lichnerowicz equation for (γ, p, t) ,

$$\Delta_\gamma \phi - \frac{1}{8} \phi + \frac{1}{12} t^2 \phi^5 - \frac{1}{8} (p \cdot p) \mu^{-2} \phi^{-7} = 0$$

(here Δ_γ is the Laplacian for γ and μ is the volume form for γ). Our desired (q, π) is given by $q := \phi^4 \gamma$ and $\pi := \phi^{-4} p + \frac{2}{3} t \phi^2 \gamma^{-1}$.

and of spatial volume.¹⁹⁵

Taking this example as our model, we can introduce the notion of a *Hamiltonianization* of general relativity associated with a given parameterized geometric time for the theory. Suppose that we are given such a parameterized geometric time. Suppose further that we are able to construct a symplectic space \mathcal{I}^* whose points are ($\mathcal{D}(S)$ -equivalence classes of) tensors on the abstract instant S , and that for each value of t we are able to construct a geometrically natural isomorphism between \mathcal{I}^* and the set of instantaneous geometries corresponding to t according to our parameterized geometric time. Composing these isomorphisms with the canonical map from the space of instantaneous geometries to the reduced space of solutions gives us a one-parameter family of symplectic isomorphisms between \mathcal{I}^* and \mathcal{S}' .¹⁹⁶ This allows us to associate each geometry $[g]$ with a curve $x(t)$ in \mathcal{I}^* : for each t , $x(t)$ is the point in \mathcal{I}^* that gets mapped to $[g]$ by the isomorphism labelled by t . We call $x(t)$ the dynamical trajectory associated with $[g]$. We now consider the class of dynamical trajectories on \mathcal{I}^* that arise in this way, and ask whether there is a (possibly time-dependent) Hamiltonian on \mathcal{I}^* that generates them in concert with the symplectic structure of \mathcal{I}^* . If there is, then the resulting (possibly time-dependent) Hamiltonian system is a Hamiltonianization of general relativity based upon the given parameterized geometric time.

As we have seen, given a parameterized geometric time for general relativity we can represent changeable quantities in the familiar way via one-parameter families of functions on the reduced space of solutions. And if we go further and introduce an associated Hamiltonianization of the theory, then we can represent time in the familiar way via a (possibly time-dependent) Hamiltonian flow on the symplectic space \mathcal{I}^* , whose points we can think of as initial data posable at different times. So these notions allow us to circumvent the problem of time by playing the same roles that a slicing played in sections 5 and 6 when we considered theories set in fixed background spacetimes.

Does the introduction of a geometric time or of an associated Hamiltonianization violate general covariance? In one sense there is no violation — for these notions are situated at the level of the reduced space of solutions, and so cannot, e.g., treat diffeomorphic solutions differently.

But it remains true that the introduction of a geometric time violates the spirit of general relativity, as the theory is generally understood today — most would like to think of special relativity as dissolving any privileged distinction between time and space and of general relativity as generalizing special relativity in a way that does nothing to reinstate such a distinction.¹⁹⁷

¹⁹⁵The spatial volume is itself a t -dependent function on \mathcal{I}^* , since the same conformal data will lead to instantaneous geometries with different volumes when supplemented by different values of t .

¹⁹⁶Strictly speaking, these isomorphisms will be merely local (as in section 6.1) if the range of values taken on by the time parameter varies from geometry to geometry.

¹⁹⁷On the other hand, many early relativistic cosmologists were happy to take the natural foliation of nonrotating dust cosmologies by surfaces orthogonal to the dust worldlines as a sign that the distinction between space and time, banished in Einstein's account of electromagnetism,

Note, however, that this is really an objection to the privileging of one geometric time over others. It seems entirely in the spirit of general relativity to think of the content of the theory as being elucidated by each of its Hamiltonianizations and as being exhausted by the set of all Hamiltonianizations (that is, if we ignore spacetimes with time translation or reflection symmetries).

Still, it is natural to ask what sort of considerations could lead us to recognize a geometric time or associated Hamiltonianization as being the *correct* one.¹⁹⁸

Classical Considerations. In the CMC Hamiltonianization sketched in example 54 above general relativity is recast as a time-dependent system. This is a bit unsettling: we are used to thinking that time-dependent Hamiltonians only arise when an open system is subject to external forces. So it is surprising to encounter a time-dependent Hamiltonian system in a fundamental context. Perhaps this is something we have to learn to live with: we are here in effect singling out one of general relativity's infinitely many variables and treating it as time — and we expect there to be all sorts of complicated nonlinear interactions between the variables of general relativity. However, some interesting special cases are known of geometric times that lead to time-independent Hamiltonianizations of general relativity.¹⁹⁹

So we cannot rule out the possibility that there may be a geometric time of wide scope that allows us to reformulate general relativity as a time-independent Hamiltonian theory with non-trivial dynamics.²⁰⁰ Clearly the construction of such a geometric time would be of the first interest: it might well seem that we had happened on the *correct* time, previously concealed from us by the unperspicuous formulations of the theory that we had been working with — much as it would have if classical mechanics had first been given a time-reparameterization invariant formulation, and it had then been discovered that a certain family of parameterizations allowed the equations to be rewritten in a much simpler form.

Quantum Considerations. The question whether to privilege one geometric time or to treat them all equally can be expected to have repercussions for quantization (which project provides the main motivation for looking for

was reinstated at the astronomical level. See [Belot, 2005, §3.2] for discussion and references.

¹⁹⁸Note that some approaches in the philosophy of time and some approaches to the interpretation of quantum mechanics would appear to require something like a preferred foliation of spacetime by instants of time.

¹⁹⁹This happens with the CMC time in the case where space has the topology of a two-torus; see [Moncrief, 1989, 2913]. It can also be achieved for general relativity coupled to a perfect fluid — in this case the conserved quantity that drives the dynamics is total baryon number; see [Moncrief, 1977] and [Moncrief and Demaret, 1980].

²⁰⁰Note that given a non-trivial time-independent Hamiltonian on \mathcal{I}^* , we can use our t -dependent isomorphisms between \mathcal{I}^* and \mathcal{S}' to construct a corresponding (possibly time-dependent) Hamiltonian on \mathcal{S}' . Unless the latter function were constant, it would generate a nontrivial flow on the reduced space of solutions. Of course, this could not be interpreted as time translation, although generated by the counterpart of the time-independent Hamiltonian generating time evolution.

a Hamiltonian formulation of general relativity with nontrivial dynamics in the first place). For one certainly does not expect that distinct Hamiltonian formulations of general relativity corresponding to distinct choices of geometric time should have equivalent quantizations — at least not if equivalent quantizations are required to be unitarily equivalent.²⁰¹

So what can we hope for? For long shots like the following. (1) Perhaps only one geometric time will lead to an empirically adequate quantum theory of gravity. (2) Perhaps there will be a natural class of geometric times (e.g., the ones that lead to time-independent Hamiltonians) that can be seen as underwriting the equivalent quantum theories (perhaps in a liberalized sense of “equivalent”).

Far more plausibly, the solution to the difficulties in quantizing general relativity will come from some other direction entirely. But hopefully it will in any case be worthwhile to be clear about the nature of the problem of time.

ACKNOWLEDGEMENTS

I would like to thank: the editors of this volume for their invitation, their sage advice, and their great patience; audiences in Montreal, New York, and Pittsburgh who sat through chaotic presentations of this material; Vincent Moncrief and Charles Torre for helpful correspondence; Richard Thomas for unhelpful correspondence; and Frank Arntzenius, Cian Dorr, Alexandre Guay, and Laura Ruetsche for helpful conversations.

BIBLIOGRAPHY

- [Abraham and Marsden, 1978] R. Abraham and J. Marsden. *Foundations of Mechanics*. Perseus Books, Cambridge, MA, 2nd edition, 1978.
- [Abraham *et al.*, 1988] R. Abraham, J. Marsden, and T. Ratiu. *Manifolds, Tensor Analysis, and Applications*. Springer-Verlag, Berlin, 2nd edition, 1988.
- [Adams *et al.*, 1985] M. Adams, T. Ratiu, and R. Schmid. The Lie group structure of diffeomorphism groups and invertible Fourier operators, with applications. In V. Kac (ed.), *Infinite-Dimensional Groups with Applications*, pages 1–69. Springer-Verlag, Berlin, 1985.
- [Ali and Engliš, 2005] T. Ali and M. Engliš. Quantization methods: A guide for physicists and analysts. *Reviews in Mathematical Physics*, 17:391–490, 2005.
- [Anderson and Duchamp, 1980] I. Anderson and T. Duchamp. On the existence of global variational principles. *American Journal of Mathematics*, 102:781–868, 1980.

²⁰¹The Stone-von Neumann theorem guarantees that unless one of them does something odd, two people setting out to quantize a Hamiltonian theory with a linear, finite-dimensional phase space will end up with unitarily equivalent quantizations. But as soon as one considers infinite-dimensional or nonlinear phase spaces the situations changes radically — for example, what look like equivalent formalisms at the classical level lead to distinct quantum theories. See [Ruetsche, Unpublished] and [Gotay, 2000] for discussion, examples, and references. See [Gotay and Demaret, 1983] for a minisuperspace cosmological model that admits competing deparameterizations that lead to physically distinct quantizations.

- [Anderson and Thompson, 1992] I. Anderson and G. Thompson. *The Inverse Problem of the Calculus of Variations for Ordinary Differential Equations*. American Mathematical Society, Providence, RI, 1992.
- [Andersson *et al.*, 1997] L. Andersson, V. Moncrief, and A. Tromba. On the global evolution problem in 2+1 gravity. *Journal of Geometry and Physics*, 23:191–205, 1997.
- [Andersson *et al.*, 1998] L. Andersson, G. Galloway, and R. Howard. The cosmological time function. *Classical and Quantum Gravity*, 15:309–322, 1998.
- [Andersson, 1987] L. Andersson. Momenta and reduction in general relativity. *Journal of Geometry and Physics*, 4:289–314, 1987.
- [Andersson, 1989] L. Andersson. Momenta and reduction in general relativity. II. The level sets. In R. Bartnik (ed.), *Conference of Mathematical Relativity*, pages 73–88. Australian National University, Canberra, 1989.
- [Andersson, 2004] L. Andersson. The global existence problem in general relativity. In P. Chruściel and H. Friedrich (eds.), *The Einstein Equations and the Large Scale Behavior of Gravitational Fields*, pages 71–120. Birkhäuser, Basel, 2004.
- [Arms, 1981] J. Arms. The structure of the solution set for the Yang–Mills equations. *Mathematical Proceedings of the Cambridge Philosophical Society*, 90:361–372, 1981.
- [Arnold *et al.*, 1997] V. I. Arnold, V. V. Kozlov, and A. I. Neistadt. *Mathematical Aspects of Classical and Celestial Mechanics*. Springer-Verlag, Berlin, 2nd edition, 1997.
- [Arnold, 1989] V. I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer-Verlag, Berlin, 2nd edition, 1989.
- [Ashtekar and Romano, 1992] A. Ashtekar and J. Romano. Spatial infinity as a boundary of spacetime. *Classical and Quantum Gravity*, 9:1069–1100, 1992.
- [Ashtekar *et al.*, 1991] A. Ashtekar, L. Bombelli, and O. Reula. The covariant phase space of asymptotically flat gravitational fields. In M. Francaviglia (ed.), *Mechanics, Analysis, and Geometry: 200 Years After Lagrange*, pages 417–450. Elsevier, Amsterdam, 1991.
- [Barbot and Zeghib, 2004] T. Barbot and A. Zeghib. Group actions on Lorentz spaces, mathematical aspects: A survey. In P. Chruściel and H. Friedrich (eds.), *The Einstein Equation and the Large Scale Behavior of Gravitational Fields*, pages 401–439. Birkhäuser, Basel, 2004.
- [Barbot, 2005] T. Barbot. Flat globally hyperbolic spacetimes. *Journal of Geometry and Physics*, 53:123–165, 2005.
- [Bartnik and Isenberg, 2004] R. Bartnik and J. Isenberg. The constraint equations. In P. Chruściel and H. Friedrich (eds.), *The Einstein Equations and the Large Scale Behavior of Gravitational Fields*, pages 1–38. Birkhäuser, Basel, 2004.
- [Bartnik *et al.*, 1990] R. Bartnik, P. Chruściel, and N. Ó Murchadha. On maximal surfaces in asymptotically flat spacetimes. *Communications in Mathematical Physics*, 130:95–109, 1990.
- [Bartnik, 1988] R. Bartnik. Remarks on cosmological spacetimes and constant mean curvature surfaces. *Communications in Mathematical Physics*, 117:615–624, 1988.
- [Beig and Ó Murchadha, 1987] R. Beig and N. Ó Murchadha. The Poincaré group as the symmetry group of canonical general relativity. *Annals of Physics*, 174:463–498, 1987.
- [Beig and Schmidt, 1982] R. Beig and B. Schmidt. Einstein’s equations near spatial infinity. *Communications in Mathematical Physics*, 87:65–80, 1982.
- [Beig, 1994] R. Beig. The classical theory of canonical general relativity. In J. Ehlers and H. Friedrich (eds.), *Canonical Gravity: From Classical to Quantum*, 59–80. Springer-Verlag, Berlin, 1994.
- [Belot and Earman, 2001] G. Belot and J. Earman. Pre-Socratic quantum gravity. In C. Callender and N. Huggett (eds.), *Physics Meets Philosophy and the Planck Scale*, pages 213–255. Cambridge University Press, Cambridge, 2001.
- [Belot, 2003] G. Belot. Symmetry and gauge freedom. *Studies in History and Philosophy of Modern Physics*, 34:189–225, 2003.
- [Belot, 2005] G. Belot. Dust, time, and symmetry. *British Journal for the Philosophy of Science*, 55:255–291, 2005.
- [Benedetti and Guadagnini, 2001] R. Benedetti and E. Guadagnini. Cosmological time in (2+1)-gravity. *Nuclear Physics B*, 613:330–352, 2001.
- [Bluman, 2005] G. Bluman. Connections between symmetries and conservation laws. *Symmetry, Integrability and Geometry: Methods and Applications*, 1:011, 2005.
- [Brading and Castellani, this volume] K. Brading and E. Castellani. Symmetries and invariances in classical physics. Chapter 13, this volume.

- [Budic *et al.*, 1978] R. Budic, J. Isenberg, L. Lindblom, and P. Yasskin. On the determination of Cauchy surfaces from intrinsic properties. *Communications in Mathematical Physics*, 61:87–95, 1978.
- [Butterfield and Isham, 2000] J. Butterfield and C. Isham. On the emergence of time in quantum gravity. In J. Butterfield (ed.), *The Arguments of Time*, pages 111–168. Oxford University Press, Oxford, 2000.
- [Butterfield, this volume] J. Butterfield. On symplectic reduction in mechanics. Chapter 1, this volume.
- [Butterfield, Unpublished] J. Butterfield. Against *Pointillisme* about mechanics. Unpublished.
- [Cannas da Silva, 2006] A. Cannas da Silva. Symplectic geometry. In F. Dillen and L. Verstraeten, eds. *Handbook of Differential Geometry, Volume II*, pp. 79–188, North-Holland, 2006.
- [Cantor, 1979] M. Cantor. Some problems of global analysis on asymptotically simple manifolds. *Compositio Mathematica*, 38:3–35, 1979.
- [Choquet-Bruhat and DeWitt-Morrette, 1989] Y. Choquet-Bruhat and C. DeWitt-Morrette. *Analysis, Manifolds, and Physics. Part II: 92 Applications*. North-Holland, Amsterdam, 1989.
- [Choquet-Bruhat *et al.*, 1977] Y. Choquet-Bruhat, C. DeWitt-Morrette, and M. Dillard-Bleick. *Analysis, Manifolds, and Physics*. North-Holland, Amsterdam, 1977.
- [Chruściel *et al.*, 2005] J. Chruściel, J. Isenberg and D. Pollack. Initial data engineering. *Communications in Mathematical Physics*, 257:29–42, 2005.
- [Crampin and Prince, 1988] M. Crampin and G. E. Prince. Alternative Lagrangians for spherically symmetric potentials. *Journal of Mathematical Physics*, 29:1551–1554, 1988.
- [Crnković and Witten, 1987] Č. Crnković and E. Witten. Covariant description of canonical formalism in geometrical theories. In S. Hawking and W. Israel (eds.), *Three Hundred Years of Gravitation*, pages 676–684. Cambridge University Press, Cambridge, 1987.
- [Cushman and Bates, 1997] R. Cushman and L. Bates. *Global Aspects of Classical Integrable Systems*. Birkhäuser, Basel, 1997.
- [Deligne and Freed, 1999] P. Deligne and D. Freed. Classical field theory. In P. Deligne, P. Etingof, D. Freed, L. Jeffrey, D. Kazhdan, J. Morgan, D. Morrison, and E. Witten (eds.), *Quantum Fields and Strings: A Course for Mathematicians*, volume 1, pages 137–225. American Mathematical Society, Providence, RI, 1999.
- [Diacu and Holmes, 1996] F. Diacu and P. Holmes. *Celestial Encounters*. Princeton University Press, Princeton, 1996.
- [Diacu, 1992] F. Diacu. *Singularities of the N-Body Problem*. Les Publications CRM, Montréal, 1992.
- [Diacu, 2002] F. Diacu. Singularities of the N-body problem. In H. Cabral and F. Diacu (eds.), *Classical and Celestial Mechanics: The Recife Lectures*, pages 35–62. Princeton University Press, Princeton, 2002.
- [Dirac, 2001] P. A. M. Dirac. *Lectures on Quantum Mechanics*. Dover, New York, 2001.
- [Dubrovín *et al.*, 1992] B. A. Dubrovín, A. T. Fomenko, and S. P. Novikov. *Modern Geometry — Methods and Applications. Part I. The Geometry of Surfaces, Transformation Groups, and Fields*. Springer-Verlag, Berlin, 2nd edition, 1992.
- [Duval *et al.*, 1990] C. Duval, J. Elhadad, M. Gotay, and G. Tuynman. Nonunimodularity and the quantization of the pseudo-rigid body. In J. Harnad and J. Marsden (eds.), *Hamiltonian Systems, Transformation Groups, and Spectral Transform Methods*, pages 149–160. Les Publications CRM, Montréal, 1990.
- [Eardley and Smarr, 1979] D. Eardley and L. Smarr. Time functions in numerical relativity: Marginally bound dust collapse. *Physical Review D*, 19:2239–2259, 1979.
- [Earman, 1986] J. Earman. *A Primer on Determinism*. Reidel, Dordrecht, 1986.
- [Earman, 2002] J. Earman. Thoroughly modern McTaggart, or what McTaggart would have said if he had read the general theory of relativity. *Philosophers' Imprint*, 2(3):1–28, 2002.
- [Earman, 2003] J. Earman. Tracking down gauge. In K. Brading and E. Castellani (eds.), *Symmetries in Physics*, pages 140–162. Cambridge University Press, Cambridge, 2003.
- [Earman, this volume] J. Earman. Aspects of determinism in modern physics. Chapter 15, this volume.
- [Earman, 2006] J. Earman. Two challenges to the requirement of substantive general covariance. *Synthese*, 142:443–468, 2006.
- [Eichhorn, 1993] J. Eichhorn. The manifold structure of maps between open manifolds. *Annals of Global Analysis and Geometry*, 11:253–300, 1993.

- [Fischer and Moncrief, 1996] A. Fischer and V. Moncrief. A method of reduction of Einstein's equations of evolution and a natural symplectic structure on the space of gravitational degrees of freedom. *General Relativity and Gravitation*, 28:207–219, 1996.
- [Fischer and Moncrief, 1997] A. Fischer and V. Moncrief. Hamiltonian reduction of Einstein's equations of general relativity. *Nuclear Physics B (Proceedings Supplement)*, 57:142–161, 1997.
- [Fischer and Moncrief, Unpublished] A. Fischer and V. Moncrief. Nonautonomous dynamical systems and the phase portrait of the reduced einstein equations. Unpublished.
- [Frauendiener and Sparling, 1992] J. Frauendiener and G. Sparling. On the symplectic formalism for general relativity. *Proceedings of the Royal Society. London. Series A*, 436:141–153, 1992.
- [Friedman, 1983] M. Friedman. *Foundations of Spacetime Theories*. Princeton University Press, Princeton, 1983.
- [Gerver, 2003] J. Gerver. Noncollision singularities: Do four bodies suffice? *Experimental Mathematics*, 12:187–198, 2003.
- [Goldstein, 1953] H. Goldstein. *Classical Mechanics*. Addison–Wesley, Cambridge, MA, 1953.
- [Gotay and Demaret, 1983] M. Gotay and J. Demaret. Quantum cosmological singularities. *Physical Review D*, 28:2404–2413, 1983.
- [Gotay and Isenberg, 1992] M. Gotay and J. Isenberg. La symplectification de la science. *La Gazette des Mathématiciens*, 54:59–79, 1992.
- [Gotay and Nester, 1980] M. Gotay and J. Nester. Generalized constraint algorithm and special presymplectic manifolds. In G. Kaiser and J. Marsden (eds.), *Geometric Methods in Mathematical Physics*, pages 79–104. Springer-Verlag, Berlin, 1980.
- [Gotay et al., 1978] M. Gotay, J. Nester, and G. Hinds. Presymplectic manifolds and the Dirac–Bergmann theory of constraints. *Journal of Mathematical Physics*, 19:2388–2399, 1978.
- [Gotay, 1983] M. Gotay. On the validity of Dirac's conjecture regarding first-class secondary constraints. *Journal of Physics A*, 16:L141–L145, 1983.
- [Gotay, 2000] M. Gotay. Obstructions to quantization. In *Mechanics: From Theory to Computation*, pages 171–216. Springer-Verlag, Berlin, 2000.
- [Guillemin and Sternberg, 1990] V. Guillemin and S. Sternberg. *Variations on a Theme by Kepler*. American Mathematical Society, Providence, RI, 1990.
- [Healey, 2002] R. Healey. Can physics coherently deny the reality of time? In C. Callender (ed.), *Time, Reality, & Experience*, pages 293–316. Cambridge University Press, Cambridge, 2002.
- [Healey, Unpublished] R. Healey. Gauging what's real. Unpublished.
- [Henneaux and Shepley, 1982] M. Henneaux and L. Shepley. Lagrangians for spherically symmetric potentials. *Journal of Mathematical Physics*, 11:2101–2107, 1982.
- [Henneaux and Teitelboim, 1992] M. Henneaux and C. Teitelboim. *Quantization of Gauge Systems*. Princeton University Press, Princeton, 1992.
- [Hojman and Shepley, 1991] S. Hojman and L. Shepley. No Lagrangian? No quantization! *Journal of Mathematical Physics*, 32:142–146, 1991.
- [Iglesias, 1998] P. Iglesias. Les origines du calcul symplectique chez Lagrange. *L'Enseignement Mathématique*, 44:257–277, 1998.
- [Isenberg and Marsden, 1982] J. Isenberg and J. Marsden. A slice theorem for the space of solutions of Einstein's equations. *Physics Reports*, 89:179–222, 1982.
- [Isenberg and Rendall, 1998] J. Isenberg and A. Rendall. Cosmological spacetimes not covered by a constant mean curvature slicing. *Classical and Quantum Gravity*, 15:3679–3688, 1998.
- [Isham, 1993] C. Isham. Canonical quantum gravity and the problem on time. In L. A. Ibort and M. A. Rodriguez (eds.), *Integrable Systems, Quantum Groups, and Quantum Field Theories*, pages 157–288. Kluwer, Dordrecht, 1993.
- [Kay, 1980] B. Kay. Linear spin-zero quantum fields in external gravitational and scalar fields. II. Generally covariant perturbation theory. *Communications in Mathematical Physics*, 71:29–46, 1980.
- [King and Ellis, 1973] A. R. King and G. F. R. Ellis. Tilted homogeneous cosmological models. *Communications in Mathematical Physics*, 31:209–242, 1973.
- [Kondracki and Rogulski, 1986] W. Kondracki and J. Rogulski. On the stratification of the orbit space for the action of automorphisms on connections. *Dissertationes Mathematicae*, CCL:1–62, 1986.

- [Kuchař, 1986] K. Kuchař. Hamiltonian dynamics of gauge theories. *Physical Review D*, 34:3031–3043, 1986.
- [Kuchař, 1992] K. Kuchař. Time and interpretations of quantum gravity. In G. Kunstatter, D. Vincent, and J. Williams (eds.), *Proceedings of the 4th Canadian Conference on General Relativity and Astrophysics*, pages 211–314. World Scientific, Singapore, 1992.
- [Kuchař, 1993] K. Kuchař. Canonical quantum gravity. In R. Gleiser, C. Kozameh, and O. Moreschi (eds.), *General Relativity and Gravitation 1992*, pages 119–150. Institute of Physics Publishing, Philadelphia, 1993.
- [Lanczos, 1986] C. Lanczos. *Variational Principles of Mechanics*. Dover, New York, 4th edition, 1986.
- [Landsman, 1998] N. P. Landsman. Quantization of singular systems and incomplete motions. In M. Rainer and H.-J. Schmidt (eds.), *Current Topics in Mathematical Cosmology*, pages 256–263. World Scientific, Singapore, 1998.
- [Landsman, this volume] N. P. Landsman. Between classical and quantum. Chapter 5, this volume.
- [Lang, 1999] S. Lang. *Fundamentals of Differential Geometry*. Springer-Verlag, Berlin, 1999.
- [Lee and Wald, 1990] J. Lee and R. Wald. Local symmetries and constraints. *Journal of Mathematical Physics*, 31:725–743, 1990.
- [Lee, 2000] J. Lee. *Introduction to Topological Manifolds*. Springer-Verlag, Berlin, 2000.
- [Lee, 2003] J. Lee. *Introduction to Smooth Manifolds*. Springer-Verlag, Berlin, 2003.
- [Lewis, 1986] D. Lewis. *Philosophical Papers*, volume II. Oxford University Press, Oxford, 1986.
- [Lewis, 1999] D. Lewis. Humean Supervenience debugged. In *Papers in Metaphysics and Epistemology*, pages 224–247. Cambridge University Press, Cambridge, 1999.
- [Loll, 1992] R. Loll. Canonical and BRST-quantization of constrained systems. In M. Gotay, J. Marsden, and V. Moncrief (eds.), *Mathematical Aspects of Classical Field Theory*, pages 503–530. American Mathematical Society, Providence, RI, 1992.
- [Malament, this volume] D. Malament. Classical relativity theory. Chapter 3, this volume.
- [Marco, 1990a] J.-P. Marco. Chirurgie et régularisation des variétés de mouvements. *Comptes Rendus des Séances de l'Académie des Sciences*, 310:125–127, 1990.
- [Marco, 1990b] J.-P. Marco. Variétés de mouvements et régularisation des systèmes dynamiques. *Comptes Rendus des Séances de l'Académie des Sciences*, 310:65–67, 1990.
- [Marsden and Ratiu, 1994] J. Marsden and T. Ratiu. *Introduction to Mechanics and Symmetry*. Springer-Verlag, Berlin, 1994.
- [Marsden and Weinstein, 1982] J. Marsden and A. Weinstein. The Hamiltonian structure of the Maxwell–Vlasov equations. *Physica D*, 4:394–406, 1982.
- [Marsden *et al.*, 1972] J. Marsden, D. Ebin and A. Fischer. Diffeomorphism groups, hydrodynamics and relativity. In R. Vanstone (ed.), *Proceedings of the Thirteenth Biennial Seminar of the Canadian Mathematical Congress on Differential Geometry and Applications*, volume 1, pages 135–279. Canadian Mathematical Congress, Montreal, 1972.
- [Marsden, 1981] J. Marsden. *Lectures on Geometric Methods in Mathematical Physics*. SIAM, Philadelphia, 1981.
- [Maudlin, 2002] T. Maudlin. Thoroughly muddled McTaggart, or how to abuse gauge freedom to generate metaphysical monstrosities. *Philosophers' Imprint*, 2(4):1–19, 2002.
- [Maudlin, Unpublished] T. Maudlin. Suggestions from physics for deep metaphysics. Unpublished.
- [McGehee, 1975] R. McGehee. Triple collisions in gravitational systems. In J. Moser (ed.), *Dynamical Systems, Theory and Application*, pages 550–572. Springer-Verlag, Berlin, 1975.
- [Milnor, 1984] J. Milnor. Remarks on infinite-dimensional Lie groups. In B. DeWitt and R. Stora (eds.), *Relativity, Groups, and Topology II*, 1007–1057. North-Holland, Amsterdam, 1984.
- [Moncrief and Demaret, 1980] V. Moncrief and J. Demaret. Hamiltonian formalism for perfect fluids in general relativity. *Physical Review D*, 21:2785–2793, 1980.
- [Moncrief, 1977] V. Moncrief. Hamiltonian formalism for relativistic perfect fluids. *Physical Review D*, 16:1702–1705, 1977.
- [Moncrief, 1989] V. Moncrief. Reduction of the Einstein equations in 2+1 dimensions to a Hamiltonian system over Teichmüller space. *Journal of Mathematical Physics*, 30:2907–2914, 1989.
- [Moncrief, 1992] V. Moncrief. Boost symmetries in spatially compact spacetimes with a cosmological constant. *Classical and Quantum Gravity*, 9:2515–2520, 1992.

- [Norton, 1995] J. Norton. Did Einstein stumble: The debate over general covariance. *Erkenntnis*, 42:223–245, 1995.
- [Nutku and Pavlov, 2002] Y. Nutku and M. Pavlov. Multi-Lagrangians for integrable systems. *Journal of Mathematical Physics*, 43:1441–1459, 2002.
- [Olver, 1993] P. Olver. *Applications of Lie Groups to Differential Equations*. Springer-Verlag, Berlin, 2nd edition, 1993.
- [Ortega and Ratiu, 2004] J.-P. Ortega and T. Ratiu. *Momentum Maps and Hamiltonian Reduction*. Birkhäuser, Basel, 2004.
- [Pflaum, 2001] M. Pflaum. *Analytic and Geometric Study of Stratified Spaces*. Springer-Verlag, Berlin, 2001.
- [Prince, 2000] G. E. Prince. The inverse problem in the calculus of variations and its ramifications. In P. Vassiliou and I. Lisle (eds.), *Geometric Approaches to Differential Equations*, pages 171–200. Cambridge University Press, Cambridge, 2000.
- [Redhead, 2003] M. Redhead. The interpretation of gauge symmetry. In K. Brading and E. Castellani (eds.), *Symmetries in Physics*, pages 124–139. Cambridge University Press, Cambridge, 2003.
- [Reichenbach, 1991] H. Reichenbach. *The Direction of Time*. University of California Press, Berkeley, 1991.
- [Rendall, 1996] A. Rendall. Constant mean curvature foliations in cosmological spacetimes. *Helvetica Physica Acta*, 69:490–500, 1996.
- [Rovelli, 1991] C. Rovelli. Time in quantum gravity: A hypothesis. *Physical Review D*, 43:442–456, 1991.
- [Rovelli, this volume] C. Rovelli. Quantum gravity. Chapter 12, this volume.
- [Ruetsche, Unpublished] L. Ruetsche. Quantizing. Unpublished.
- [Rynasiewicz, 1999] R. Rynasiewicz. Kretschmann’s analysis of general covariance. In H. Goenner, J. Renn, J. Ritter, and T. Sauer (eds.), *The Expanding Worlds of General Relativity*, 431–462. Birkhäuser, Basel, 1999.
- [Saari and Xia, 1995] D. Saari and Z. Xia. Off to infinity in finite time. *Notices of the American Mathematical Society*, 42:538–546, 1995.
- [Sachs and Wu, 1977] R. Sachs and H.-H. Wu. *General Relativity for Mathematicians*. Springer-Verlag, Berlin, 1977.
- [Schmid, 1987] R. Schmid. *Infinite Dimensional Hamiltonian Systems*. Bibliopolis, Naples, 1987.
- [Singer, 2001] S. Singer. *Symmetry in Mechanics: A Gentle, Modern Introduction*. Birkhäuser, Basel, 2001.
- [Sorkin, 2002] R. Sorkin. An example relevant to the Kretschmann–Einstein debate. *Modern Physics Letters A*, 17:695–700, 2002.
- [Souriau, 1982] J.-M. Souriau. Géométrie globale du problème à deux corps. *Atti della Accademia delle Scienze di Torino. Classe di Scienze Fisiche, Matematiche e Naturali*, 117(Supplement 1):369–418, 1982.
- [Souriau, 1986] J.-M. Souriau. La structure symplectique décrite par lagrange en 1811. *Mathématiques et Sciences Humaines*, 94:45–54, 1986.
- [Souriau, 1997] J.-M. Souriau. *Structure of Dynamical Systems*. Birkhäuser, Basel, 1997.
- [Sudbery, 1986] A. Sudbery. A vector Lagrangian for the electromagnetic field. *Journal of Physics A*, 19:L33–L36, 1986.
- [Tarkovsky, 1991] A. Tarkovsky. *Time Within Time: The Diaries, 1970–1986*. Seagull, Calcutta, 1991.
- [Torre and Anderson, 1996] C. Torre and I. Anderson. Classification of local generalized symmetries for the vacuum Einstein equations. *Communications in Mathematical Physics*, 176:479–539, 1996.
- [Torre, 1992] C. Torre. Covariant phase space formulation of parameterized field theories. *Journal of Mathematical Physics*, 33:3802–3811, 1992.
- [Torre, Unpublished] C. Torre. Local cohomology in field theory (with applications to the Einstein equations). Unpublished.
- [Tuynman, 1992] G. Tuynman. What are the rules of the game called BRST? In M. Gotay, J. Marsden, and V. Moncrief (eds.), *Mathematical Aspects of Classical Field Theory*, pages 625–633. American Mathematical Society, Providence, RI, 1992.
- [van Brunt, 2004] B. van Brunt. *The Calculus of Variations*. Springer-Verlag, Berlin, 2004.
- [Wald, 1984] R. Wald. *General Relativity*. University of Chicago Press, Chicago, 1984.

- [Wald, 1994] R. Wald. *Quantum Field Theory on Curved Spacetime*. University of Chicago Press, Chicago, 1994.
- [Wallace, 2003] D. Wallace. Time-dependent symmetries: The link between gauge symmetries and indeterminism. In K. Brading and E. Castellani (eds.), *Symmetries in Physics*, pages 163–173. Cambridge University Press, Cambridge, 2003.
- [Weinstein, 1981] A. Weinstein. Symplectic geometry. *Bulletin of the American Mathematical Society*, 5:1–13, 1981.
- [Woodhouse, 1991] N. M. J. Woodhouse. *Geometric Quantization*. Oxford University Press, Oxford, 2nd edition, 1991.
- [Zuckerman, 1987] G. Zuckerman. Action principles and global geometry. In S.-T. Yau (ed.), *Mathematical Aspects of String Theory*, pages 259–284. World Scientific, 1987.

This page intentionally left blank

CLASSICAL RELATIVITY THEORY

David B. Malament

1 INTRODUCTION

The essay that follows is divided into two parts. In the first, I give a brief account of the structure of classical relativity theory.¹ In the second, I discuss three special topics.

My account in the first part (section 2) is limited in several respects. I do not discuss the historical development of classical relativity theory, nor the evidence we have for it. I do not treat “special relativity” as a theory in its own right that is superseded by “general relativity”. And I do not describe known exact solutions to Einstein’s equation. (This list could be continued at great length.²)

Instead, I limit myself to a few fundamental ideas, and present them as clearly and precisely as I can. The account presupposes a good understanding of basic differential geometry, and at least passing acquaintance with relativity theory itself.³

In section 3, I first consider the status of the relative simultaneity relation in the context of Minkowski spacetime. At issue is whether the standard relation, the one picked out by Einstein’s “definition” of simultaneity, is conventional in character, or is rather in some significant sense forced on us. Then I describe the “geometrized” version of Newtonian gravitation theory (also known as Newton-Cartan theory). It is included here because it helps to clarify what is and is not distinctive about classical relativity theory. Finally, I consider to what extent the global geometric structure of spacetime can be recovered from its “causal structure”.⁴

¹I speak of “classical” relativity theory because considerations involving quantum mechanics will play no role. In particular, there will be no discussion of quantum field theory in curved spacetime, or of attempts to formulate a quantum theory of gravitation. (For the latter, see Rovelli (this volume, chapter 12).)

²Two important topics that I do not consider figure centrally in other contributions to this volume, namely the initial value formulation of relativity theory (Earman, chapter 15), and the Hamiltonian formulation of relativity theory (Belot, chapter 2).

³A review of the needed differential geometry (and “abstract-index notation” that I use) can be found, for example, in Wald [1984] and Malament [unpublished]. (Some topics are also reviewed in sections 3.1 and 3.2 of Butterfield (this volume, chapter 1).) In preparing part 1, I have drawn heavily on a number of sources. At the top of the list are Geroch [unpublished], Hawking and Ellis [1972], O’Neill [1983], Sachs and Wu [1977a; 1977b], and Wald [1984].

⁴Further discussion of the foundations of classical relativity theory, from a slightly different point of view, can be found in Rovelli (this volume, chapter 12).

2 THE STRUCTURE OF RELATIVITY THEORY

2.1 Relativistic Spacetimes

Relativity theory determines a class of geometric models for the spacetime structure of our universe (and subregions thereof such as, for example, our solar system). Each represents a possible world (or world-region) compatible with the constraints of the theory. It is convenient to describe these models in stages. We start by characterizing a broad class of “relativistic spacetimes”, and discussing their interpretation. Later we introduce further restrictions involving global spacetime structure and Einstein’s equation.

We take a *relativistic spacetime* to be a pair (M, g_{ab}) , where M is a smooth, connected, four-dimensional manifold, and g_{ab} is a smooth, semi-Riemannian metric on M of Lorentz signature $(1, 3)$.⁵

We interpret M as the manifold of point “events” in the world.⁶ The interpretation of g_{ab} is given by a network of interconnected physical principles. We list three in this section that are relatively simple in character because they make reference only to point particles and light rays. (These objects alone suffice to determine the metric, at least up to a constant.) In the next section, we list a fourth that concerns the behavior of (ideal) clocks. Still other principles involving generic matter fields will come up later.

We begin by reviewing a few definitions. In what follows, let (M, g_{ab}) be a fixed relativistic spacetime, and let ∇_a be the derivative operator on M determined by g_{ab} , i.e. the unique (torsion-free) derivative operator on M satisfying the compatibility condition $\nabla_a g_{bc} = 0$.

Given a point p in M , and a vector η^a in the tangent space M_p at p , we say η^a is:

<i>timelike</i>	if	$\eta^a \eta_a$	$>$	0
<i>null (or lightlike)</i>	if	$\eta^a \eta_a$	$=$	0
<i>causal</i>	if	$\eta^a \eta_a$	\geq	0
<i>spacelike</i>	if	$\eta^a \eta_a$	$<$	0

⁵The stated signature condition is equivalent to the requirement that, at every point p in M , the tangent space M_p have a basis ξ^1, \dots, ξ^4 such that, for all i and j in $\{1, 2, 3, 4\}$, $g_{ab} \xi^a \xi^b = 0$ if $i \neq j$, and

$$g_{ab} \xi^a \xi^b = \begin{cases} +1 & \text{if } i = 1 \\ -1 & \text{if } i = 2, 3, 4. \end{cases}$$

(Here we are using the abstract-index notation. ‘ a ’ is an abstract index, while ‘ i ’ and ‘ j ’ are normal counting indices.) It follows that given any vectors $\eta^a = \sum_{i=1}^4 k^i \xi^i$, and $\rho^a = \sum_{j=1}^4 l^j \xi^j$ at p ,

$$\begin{aligned} g_{ab} \eta^a \rho^b &= k^1 l^1 - k^2 l^2 - k^3 l^3 - k^4 l^4 \\ g_{ab} \eta^a \eta^b &= k^1 k^1 - k^2 k^2 - k^3 k^3 - k^4 k^4. \end{aligned}$$

In what follows, we will often use the standard convention for lowering (abstract) indices with the metric g_{ab} , and raising them with the inverse metric g^{ab} . So, for example, we will write $\eta_a \rho^a$ or $\eta^a \rho_a$ instead of $g_{ab} \eta^a \rho^b$.

⁶We use ‘event’ as a neutral term here and intend no special significance. Some might prefer to speak of “equivalence classes of coincident point events”, or “point event locations”, or something along those lines.

In this way, g_{ab} determines a “null-cone structure” in the tangent space at every point of M . Null vectors form the boundary of the cone. Timelike vectors form its interior. Spacelike vectors fall outside the cone. Causal vectors are those that are either timelike or null. This classification extends naturally to *curves*. We take these to be smooth maps of the form $\gamma: I \rightarrow M$ where $I \subseteq \mathbb{R}$ is a (possibly infinite, not necessarily open) interval.⁷ γ qualifies as *timelike* (respectively *null*, *causal*, *spacelike*) if its tangent vector field $\vec{\gamma}$ is of this character at every point.

A curve $\gamma_2: I_2 \rightarrow M$ is called an (*orientation preserving*) *reparametrization* of the curve $\gamma_1: I_1 \rightarrow M$ if there is a smooth map $\tau: I_2 \rightarrow I_1$ of I_2 onto I_1 , with positive derivative everywhere, such that $\gamma_2 = (\gamma_1 \circ \tau)$. The property of being timelike, null, etc. is preserved under reparametrization.⁸ So there is a clear sense in which our classification also extends to *images* of curves.⁹

A curve $\gamma: I \rightarrow M$ is said to be a *geodesic* (*with respect to g_{ab}*) if its tangent field ξ^a satisfies the condition: $\xi^n \nabla_n \xi^a = 0$. The property of being a geodesic is not, in general, preserved under reparametrization. So it does not transfer to curve images. But, of course, the related property of being a *geodesic up to reparametrization* does carry over. (The latter holds of a curve if it can be reparametrized so as to be a geodesic.)

Now we can state the first three interpretive principles. For all curves $\gamma: I \rightarrow M$,

- C1 γ is timelike iff its image $\gamma[I]$ could be the worldline of a massive point particle (i.e. a particle with positive mass);¹⁰
- C2 γ can be reparametrized so as to be a null geodesic iff $\gamma[I]$ could be the trajectory of a light ray;¹¹
- P1 γ can be reparametrized so as to be a timelike geodesic iff $\gamma[I]$ could be the worldline of a *free*¹² massive point particle.

In each case, a statement about geometric structure (on the left) is correlated with a statement about the behavior of particles or light rays (on the right).

Several comments and qualifications are called for. First, we are here working within the framework of relativity as traditionally understood, and ignoring speculations about

⁷If I is not an open set, we can understand smoothness to mean that there is an open interval $\bar{I} \subseteq \mathbb{R}$, with $I \subset \bar{I}$, and a smooth map $\bar{\gamma}: \bar{I} \rightarrow M$, such that $\bar{\gamma}(s) = \gamma(s)$ for all $s \in I$.

⁸This follows from the fact that, in the case just described, $\bar{\gamma}_2 = \frac{d\tau}{ds} \bar{\gamma}_1$, with $\frac{d\tau}{ds} > 0$.

⁹The difference between curves and curve images, i.e. between maps $\gamma: I \rightarrow M$ and sets $\gamma[I]$, matters. We take worldlines to be instances of the latter, i.e. construe them as point sets rather than parametrized point sets.

¹⁰We will later discuss the concept of mass in relativity theory. For the moment, we take it to be just a primitive attribute of particles.

¹¹For certain purposes, even within classical relativity theory, it is useful to think of light as constituted by streams of “photons”, and take the right side condition here to be “ $\gamma[I]$ could be the worldline of a photon”. The latter formulation makes C2 look more like C1 and P1, and draws attention to the fact that the distinction between massive particles and mass 0 particles (such as photons) has direct significance in terms of relativistic spacetime structure.

¹²“Free particles” here must be understood as ones that do not experience any forces (except “gravity”). It is one of the fundamental principles of relativity theory that gravity arises as a manifestation of spacetime curvature, not as an external force that deflects particles from their natural, straight (geodesic) trajectories. We will discuss this matter further in section 2.4.

the possibility of particles (“tachyons”) that travel faster than light. (Their worldlines would come out as images of spacelike curves.) Second, we have built in the requirement that “curves” be smooth. So, depending on how one models collisions of point particles, one might want to restrict attention here to particles that do not experience collisions.

Third, the assertions require qualification because the status of “point particles” in relativity theory is a delicate matter. At issue is whether one treats a particle’s own mass-energy as a source for the surrounding metric field g_{ab} — in addition to other sources that may happen to be present. (Here we anticipate our discussion of Einstein’s equation.) If one does, then the curvature associated with g_{ab} may blow up as one approaches the particle’s worldline. And in this case one *cannot* represent the worldline as the image of a curve in M , at least not without giving up the requirement that g_{ab} be a smooth field on M . For this reason, a more careful formulation of the principles would restrict attention to “test particles”, i.e. ones whose own mass-energy is negligible and may be ignored for the purposes at hand.

Fourth, the modal character of the assertions (i.e. the reference to possibility) is essential. It is simply not true, to take the case of C1, that all timelike curve images *are*, in fact, the worldlines of massive particles. The claim is that, as least so far as the laws of relativity theory are concerned, they *could* be. Of course, judgments concerning what could be the case depend on what conditions are held fixed in the background. The claim that a particular curve image could be the worldline of a massive point particle must be understood to mean that it could so long as there are, for example, no barriers in the way. Similarly, in C2 there is an implicit qualification. We are considering what trajectories are available to light rays when no intervening material media are present, i.e. when we are dealing with light rays *in vacua*.

Though these four concerns are important and raise interesting questions about the role of idealization and modality in the formulation of physical theory, they have little to do with relativity theory as such. Similar difficulties arise when one attempts to formulate corresponding principles within the framework of Newtonian gravitation theory.

It follows from the cited interpretive principles that the metric g_{ab} is determined (up to a constant) by the behavior of point particles and light rays. We make this claim precise in a pair of propositions about “conformal structure” and “projective structure”.

Let \bar{g}_{ab} be a second smooth metric of Lorentz signature on M . We say that \bar{g}_{ab} is *conformally equivalent* to g_{ab} if there is a smooth map $\Omega : M \rightarrow \mathbb{R}$ on M such that $\bar{g}_{ab} = \Omega^2 g_{ab}$. (Ω is called a *conformal factor*. It certainly need not be constant.) Clearly, if \bar{g}_{ab} and g_{ab} are conformally equivalent, then they agree in their classification of vectors and curves as timelike, null, etc.. The converse is true as well.¹³ Conformally equivalent metrics on M do not agree, in general, as to which curves on M qualify as geodesics or even just as geodesics up to reparametrization. But, it turns out, they do necessarily agree as to which *null* curves are geodesics up to reparametrization.¹⁴ And the converse is true,

¹³If the two metrics agree as to which vectors and curves belong to any one of the three categories, then they must agree on all. And in that case, they must be conformally equivalent. See Hawking and Ellis [1972, p. 61].

¹⁴This follows because the property of being the image of a null geodesic can be captured in terms of the existence or non-existence of (local) timelike and null curves connecting points in M . The relevant technical lemma can be formulated as follows.

once again.¹⁵

Putting the pieces together, we have the following proposition. Clauses (1) and (2) correspond to C1 and C2 respectively.

PROPOSITION 1. *Let \bar{g}_{ab} be a second smooth metric of Lorentz signature on M . Then the following conditions are equivalent.*

- (1) \bar{g}_{ab} and g_{ab} agree as to which curves on M are timelike.
- (2) \bar{g}_{ab} and g_{ab} agree as to which curves on M can be reparameterized so as to be null geodesics.
- (3) \bar{g}_{ab} and g_{ab} are conformally equivalent.

In this sense, the spacetime metric g_{ab} is determined up to a conformal factor, independently, by the set of possible worldlines of massive point particles, and by the set of possible trajectories of light rays.

Next we turn to projective structure. Let $\bar{\nabla}_a$ be a second derivative operator on M . We say that $\bar{\nabla}_a$ and ∇_a are *projectively equivalent* if they agree as to which curves are geodesics up to reparameterization (i.e. if, for all curves γ , γ can be reparametrized so as to be a geodesic with respect to $\bar{\nabla}_a$ iff it can be so reparametrized with respect to ∇_a). And if \bar{g}_{ab} is a second metric on M of Lorentz signature, we say that it is *projectively equivalent* to g_{ab} if its associated derivative operator $\bar{\nabla}_a$ is projectively equivalent to ∇_a .

It is a basic result, due to Hermann Weyl [1921], that if \bar{g}_{ab} and g_{ab} are conformally and projectively equivalent, then the conformal factor that relates them must be constant. It is convenient for our purposes, with interpretive principle P1 in mind, to cast it in a slightly altered form that makes reference only to timelike geodesics (rather than arbitrary geodesics).

PROPOSITION 2. *Let \bar{g}_{ab} be a second smooth metric on M with $\bar{g}_{ab} = \Omega^2 g_{ab}$. If \bar{g}_{ab} and g_{ab} agree as to which timelike curves can be reparameterized so as to be geodesics, then Ω is constant.*

The spacetime metric g_{ab} , we saw, is determined up to a conformal factor, independently, by the set of possible worldlines of massive point particles, and by the set of possible trajectories of light rays. The proposition now makes clear the sense in which it is fully determined (up to a constant) by those sets together with the set of possible worldlines of free massive particles.¹⁶

A curve $\gamma: I \rightarrow M$ can be reparameterized so as to be a null geodesic iff γ is null and for all $s \in I$, there is an open set $O \subseteq M$ containing $\gamma(s)$ such that, for all $s_1, s_2 \in I$, if $s_1 \leq s \leq s_2$, and if $\gamma([s_1, s_2]) \subseteq O$, then there is no timelike curve from $\gamma(s_1)$ to $\gamma(s_2)$ within O .

(Here $\gamma([s_1, s_2])$ is the image of γ as restricted to the interval $[s_1, s_2]$.) For a proof, see Hawking and Ellis [1972, p. 103].

¹⁵For if the metrics agree as to which curves are null geodesics up to reparameterization, they must agree as to which vectors at arbitrary points are null, and this, we know, implies that the metrics are conformally equivalent.

¹⁶As Weyl put it [1950, p. 103],

... it can be shown that the metrical structure of the world is already fully determined by its inertial and causal structure, that therefore mensuration need not depend on clocks and rigid bodies but that light signals and mass points moving under the influence of inertia alone will

Our characterization of relativistic spacetimes is extremely loose. Many further conditions might be imposed. For the moment, we consider just one.

(M, g_{ab}) is said to be *temporally orientable* if there exists a continuous timelike vector field τ^a on M . Suppose the condition is satisfied. Then two such fields τ^a and $\hat{\tau}^a$ on M are said to be *co-oriented* if $\tau^a \hat{\tau}_a > 0$ everywhere, i.e. if τ^a and $\hat{\tau}^a$ fall in the same lobe of the null-cone at every point of M . Co-orientation is an equivalence relation (on the set of continuous timelike vector fields on M) with two equivalence classes. A *temporal orientation* of (M, g_{ab}) is a choice of one of those two equivalence classes to count as the “future” one. Thus, a non-zero causal vector ξ^a at a point of M is said to be *future directed* or *past directed* with respect to the temporal orientation \mathcal{T} depending on whether $\tau^a \xi_a > 0$ or $\tau^a \xi_a < 0$ at the point, where τ^a is any continuous timelike vector field in \mathcal{T} . Derivatively, a causal curve $\gamma: I \rightarrow M$ is said to be *future directed* (resp. *past directed*) with respect to \mathcal{T} if its tangent vectors at every point are.

In what follows, we assume that our background spacetime (M, g_{ab}) is temporally orientable, and that a particular temporal orientation has been specified. Also, given events p and q in M , we write $p \ll q$ (resp. $p < q$) if there is a future-directed timelike (resp. causal) curve that starts at p and ends at q .¹⁷

2.2 Proper Time

So far we have discussed relativistic spacetime structure without reference to either “time” or “space”. We come to them in this section and the next.

Let $\gamma: [s_1, s_2] \rightarrow M$ be a future-directed timelike curve in M with tangent field ξ^a . We associate with it an elapsed *proper time* (relative to g_{ab}) given by

$$|\gamma| = \int_{s_1}^{s_2} (g_{ab} \xi^a \xi^b)^{\frac{1}{2}} ds.$$

This elapsed proper time is invariant under reparametrization of γ , and is just what we would otherwise describe as the length of (the image of) γ . The following is another basic principle of relativity theory.

P2 Clocks record the passage of elapsed proper time along their worldlines.

Again, a number of qualifications and comments are called for. Our formulation of C1, C2, and P1 was rough. The present formulation is that much more so. We have taken for granted that we know what “clocks” are. We have assumed that they have worldlines (rather than worldtubes). And we have overlooked the fact that ordinary clocks (e.g. the alarm clock on the nightstand) do not do well at all when subjected to extreme acceleration, tidal forces, and so forth. (Try smashing the alarm clock against the wall.) Again,

suffice.

(For more on Weyl’s “causal-inertial” method of determining the spacetime metric, see Coleman and Korté [2001, section 4.9].)

¹⁷It follows immediately that if $p \ll q$, then $p < q$. The converse does not hold, in general. But the only way the second condition can be true, without the first being true as well, is if the only future-directed causal curves from p to q are null geodesics (or reparametrizations of null geodesics). See Hawking and Ellis [1972, p. 112].

these concerns are important and raise interesting questions about the role of idealization in the formulation of physical theory. (One might construe an “ideal clock” as a point-sized test object that perfectly records the passage of proper time along its worldline, and then take P2 to assert that real clocks are, under appropriate conditions, to varying degrees of accuracy, approximately ideal.) But as with our concerns about the status of point particles, they do not have much to do with relativity theory as such. Similar ones arise when one attempts to formulate corresponding principles about clock behavior within the framework of Newtonian theory.

Now suppose that one has determined the conformal structure of spacetime, say, by using light rays. Then one can use clocks, rather than free particles, to determine the conformal factor. One has the following simple result, which should be compared with proposition 2.¹⁸

PROPOSITION 3. *Let \bar{g}_{ab} be a second smooth metric on M with $\bar{g}_{ab} = \Omega^2 g_{ab}$. Further suppose that the two metrics assign the same lengths to all timelike curves, i.e. $|\gamma|_{\bar{g}_{ab}} = |\gamma|_{g_{ab}}$ for all timelike curves $\gamma : I \rightarrow M$. Then $\Omega = 1$ everywhere. (Here $|\gamma|_{g_{ab}}$ is the length of γ relative to g_{ab} .)*

P2 gives the whole story of relativistic clock behavior (modulo the concerns noted above). In particular, it implies the path dependence of clock readings. If two clocks start at an event p , and travel along different trajectories to an event q , then, in general, they will record different elapsed times for the trip. (E.g. one will record an elapsed time of 3,806 seconds, the other 649 seconds.) This is true no matter how similar the clocks are. (We may stipulate that they came off the same assembly line.) This is the case because, as P2 asserts, the elapsed time recorded by each of the clocks is just the length of the timelike curve it traverses in getting from p to q and, in general, those lengths will be different.

Suppose we consider all future-directed timelike curves from p to q . It is natural to ask if there are any that minimize or maximize the recorded elapsed time between the events. The answer to the first question is ‘no’. Indeed, one has the following proposition.

PROPOSITION 4. *Let p and q be events in M such that $p \ll q$. Then, for all $\epsilon > 0$, there exists a future-directed timelike curve γ from p to q with $|\gamma| < \epsilon$. (But there is no such curve with length 0, since all timelike curves have non-zero length.)*

Though some work is required to give the proposition an honest proof (see O’Neill [1983, pp. 294-5]), it should seem intuitively plausible. If there is a timelike curve connecting p and q , there also exists a jointed, zig-zag null curve that connects them.

It has length 0. But we can approximate the jointed null curve arbitrarily closely with smooth timelike curves that swing back and forth. So (by the continuity of the length function), we should expect that, for all $\epsilon > 0$, there is an approximating timelike curve that has length less than ϵ . (See figure 1.)

The answer to the second question (Can one *maximize* recorded elapsed time between

¹⁸Here we not only determine the metric up to a constant, but determine the constant as well. The difference is that here, in effect, we have built in a choice of units for spacetime distance. We could obtain a more exact counterpart to proposition 2 if we worked, not with intervals of elapsed proper time, but rather with ratios of such intervals.




Figure 1. A long timelike curve from p to q and a very short one that swings back-and-forth, and approximates a broken null curve.

p and q ?) is ‘yes’ if one restricts attention to local regions of spacetime. In the case of positive definite metrics, i.e. ones with signature of form $(n, 0)$, we know, geodesics are *locally shortest* curves. The corresponding result for Lorentz metrics is that timelike geodesics are *locally longest* curves.

PROPOSITION 5. *Let $\gamma : I \rightarrow M$ be a future-directed timelike curve. Then γ can be reparametrized so as to be a geodesic iff for all $s \in I$, there exists an open set O containing $\gamma(s)$ such that, for all $s_1, s_2 \in I$ with $s_1 \leq s \leq s_2$, if the image of $\bar{\gamma} = \gamma|_{[s_1, s_2]}$ is contained in O , then $\bar{\gamma}$ (and its reparametrizations) are longer than all other timelike curves in O from $\gamma(s_1)$ to $\gamma(s_2)$. (Here $\gamma|_{[s_1, s_2]}$ is the restriction of γ to the interval $[s_1, s_2]$.)*

The proof of the proposition is very much the same as in the positive definite case. (See Hawking and Ellis [1972, p. 105].) Thus of all clocks passing locally from p to q , that one will record the greatest elapsed time that “falls freely” from p to q . To get a clock to read a smaller elapsed time than the maximal value one will have to accelerate the clock. Now acceleration requires fuel, and fuel is not free. So proposition 5 has the consequence that (locally) “saving time costs money”. And proposition 4 may be taken to imply that (locally) “with enough money one can save as much time as one wants”.

The restriction here to local regions of spacetime is essential. The connection described between clock behavior and acceleration does not, in general, hold on a global scale. In some relativistic spacetimes, one can find future-directed timelike geodesics connecting two events that have different lengths, and so clocks following the curves will record different elapsed times between the events even though *both* are in a state of free fall. Furthermore — this follows from the preceding claim by continuity considerations alone — it can be the case that of two clocks passing between the events, the one that undergoes acceleration during the trip records a *greater* elapsed time than the one that remains in a state of free fall.

The connection we have been considering between clock behavior and acceleration was once thought to be paradoxical. (I am thinking of the “clock (or twin) paradox”.)

Suppose two clocks, A and B, pass from one event to another in a suitably small region of spacetime. Further suppose A does so in a state of free fall, but B undergoes acceleration at some point along the way. Then, we know, A will record a greater elapsed time for the trip than B. This was thought paradoxical because it was believed that “relativity theory denies the possibility of distinguishing “absolutely” between free fall motion and accelerated motion”. (If we are equally well entitled to think that it is clock B that is in a state of free fall, and A that undergoes acceleration, then, by parity of reasoning, it should be B that records the greater elapsed time.) The resolution of the paradox, if one can call it that, is that relativity theory makes no such denial. The situations of A and B here are *not* symmetric. The distinction between accelerated motion and free fall makes every bit as much sense in relativity theory as it does in Newtonian physics.

In what follows, unless indication is given to the contrary, a “timelike curve” should be understood to be a future-directed timelike curve, parametrized by elapsed proper time, i.e. by arc length. In that case, the tangent field ξ^a of the curve has unit length ($\xi^a \xi_a = 1$). And if a particle happens to have the image of the curve as its worldline, then, at any point, ξ^a is called the particle’s *four-velocity* there.

2.3 Space/Time Decomposition at a Point and Particle Dynamics

Let γ be a timelike curve representing the particle O with four-velocity field ξ^a . Let p be a point on the image of γ , and let λ^a be a vector at p . There is a natural decomposition of λ^a into components parallel to, and orthogonal to, ξ^a :

$$(1) \quad \lambda^a = \underbrace{(\lambda^b \xi_b) \xi^a}_{\text{parallel to } \xi^a} + \underbrace{(\lambda^a - (\lambda^b \xi_b) \xi^a)}_{\text{orthogonal to } \xi^a}.$$

These are standardly interpreted, respectively, as the “temporal” and “spatial” components of λ^a (relative to ξ^a). In particular, the three-dimensional subspace of M_p consisting of vectors orthogonal to ξ^a is interpreted as the “infinitesimal” simultaneity slice of O at p .¹⁹ If we introduce the tangent and orthogonal projection operators

$$(2) \quad k_{ab} = \xi_a \xi_b$$

$$(3) \quad h_{ab} = g_{ab} - \xi_a \xi_b$$

then the decomposition can be expressed in the form

$$(4) \quad \lambda^a = k^a_b \lambda^b + h^a_b \lambda^b.$$

We can think of k_{ab} and h_{ab} as the relative temporal and spatial metrics determined by ξ^a . They are symmetric and satisfy

$$(5) \quad k^a_b k^b_c = k^a_c$$

$$(6) \quad h^a_b h^b_c = h^a_c.$$

¹⁹Here we simply take for granted the standard identification of “relative simultaneity” with orthogonality. We will return to consider its justification in section 3.1.

Many standard textbook assertions concerning the kinematics and dynamics of point particles can be recovered using these decomposition formulas. For example, suppose that the worldline of a second particle \bar{O} also passes through p and that its four-velocity at p is $\bar{\xi}^a$. (Since ξ^a and $\bar{\xi}^a$ are both future-directed, they are co-oriented, i.e. $(\xi^a \bar{\xi}_a) > 0$.) We compute the speed of \bar{O} as determined by O . To do so, we take the spatial magnitude of $\bar{\xi}^a$ relative to O and divide by its temporal magnitude relative to O :

$$(7) \quad v = \text{speed of } \bar{O} \text{ relative to } O = \frac{\|h^a_b \bar{\xi}^b\|}{\|k^a_b \bar{\xi}^b\|}.$$

(Given any vector μ^a , we understand $\|\mu^a\|$ to be $(\mu^a \mu_a)^{\frac{1}{2}}$ if μ^a is causal, and $(-\mu^a \mu_a)^{\frac{1}{2}}$ if it is spacelike.) From (2), (3), (5), and (6), we have

$$(8) \quad \|k^a_b \bar{\xi}^b\| = (k^a_b \bar{\xi}^b k_{ac} \bar{\xi}^c)^{\frac{1}{2}} = (k_{bc} \bar{\xi}^b \bar{\xi}^c)^{\frac{1}{2}} = (\bar{\xi}^b \xi_b)$$

and

$$(9) \quad \|h^a_b \bar{\xi}^b\| = (-h^a_b \bar{\xi}^b h_{ac} \bar{\xi}^c)^{\frac{1}{2}} = (-h_{bc} \bar{\xi}^b \bar{\xi}^c)^{\frac{1}{2}} = ((\bar{\xi}^b \xi_b)^2 - 1)^{\frac{1}{2}}.$$

So

$$(10) \quad v = \frac{((\bar{\xi}^b \xi_b)^2 - 1)^{\frac{1}{2}}}{(\bar{\xi}^b \xi_b)} < 1.$$

Thus, as measured by O , no massive particle can ever attain the maximal speed 1. (A similar calculation would show that, as determined by O , light always travels with speed 1.) For future reference, we note that (10) implies:

$$(11) \quad (\bar{\xi}^b \xi_b) = \frac{1}{\sqrt{1 - v^2}}.$$

It is a basic fact of relativistic life that there is associated with every point particle, at every event on its worldline, a *four-momentum* (or *energy-momentum*) vector P^a . In the case of a massive particle with four-velocity $\bar{\xi}^a$, P^a is proportional to $\bar{\xi}^a$, and the (positive) proportionality factor is just what we would otherwise call the *mass* (or *rest mass*) m of the particle. So we have $P^a = m \bar{\xi}^a$. In the case of a “photon” (or other mass 0 particle), no such characterization is available because its worldline is the image of a null (rather than timelike) curve. But we can still understand its four-momentum vector at the event in question to be a future-directed null vector that is tangent to its worldline there. If we think of the four-momentum vector P^a as fundamental, then we can, in both cases, recover the mass of the particle as the length of P^a : $m = (P^a P_a)^{\frac{1}{2}}$. (It is strictly positive in the first case, and 0 in the second.)

Now suppose a massive particle O has four-velocity ξ^a at an event, and another particle, either a massive particle or a photon, has four-momentum P^a there. We can recover the usual expressions for the energy and three-momentum of the second particle relative to O if we decompose P^a in terms of ξ^a . By (4) and (2), we have

$$(12) \quad P^a = (P^b \xi_b) \xi^a + h^a_b P^b.$$

The energy relative to O is the coefficient in the first term: $E = P^b \xi_b$. In the case of a massive particle where $P^a = m \xi^a$, this yields, by (11),

$$(13) \quad E = m (\bar{\xi}^b \xi_b) = \frac{m}{\sqrt{1 - v^2}}.$$

(If we had not chosen units in which $c = 1$, the numerator in the final expression would have been mc^2 and the denominator $\sqrt{1 - \frac{v^2}{c^2}}$.) The *three-momentum* relative to O is the second term in the decomposition, i.e. the component of P^a orthogonal to ξ^a : $h^a_b P^b$. In the case of a massive particle, by (9) and (11), it has magnitude

$$(14) \quad p = \|h^a_b m \bar{\xi}^b\| = m ((\bar{\xi}^b \xi_b)^2 - 1)^{\frac{1}{2}} = \frac{mv}{\sqrt{1 - v^2}}.$$

Interpretive principle P1 asserts that free particles traverse the images of timelike geodesics. It can be thought of as the relativistic version of Newton’s first law of motion. Now we consider acceleration and the relativistic version of the second law. Let $\gamma: I \rightarrow M$ be a timelike curve whose image is the worldline of a massive particle O , and let ξ^a be the four-velocity field of O . Then the *four-acceleration* (or just *acceleration*) field of O is $\xi^n \nabla_n \xi^a$, i.e. the directional derivative of ξ^a in the direction ξ^a . The four-acceleration vector is orthogonal to ξ^a . (This is clear, since $\xi^a (\xi^n \nabla_n \xi_a) = \frac{1}{2} \xi^n \nabla_n (\xi^a \xi_a) = \frac{1}{2} \xi^n \nabla_n (1) = 0$.) The magnitude $\|\xi^n \nabla_n \xi^a\|$ of the four-acceleration vector at a point is just what we would otherwise describe as the Gaussian curvature of γ there. It is a measure of the degree to which γ curves away from a straight path. (And γ is a geodesic precisely if its curvature vanishes everywhere.)

The notion of spacetime acceleration requires attention. Consider an example. Suppose you decide to end it all and jump off the Empire State Building. What would your acceleration history be like during your final moments? One is accustomed in such cases to think in terms of acceleration relative to the earth. So one would say that you undergo acceleration between the time of your jump and your calamitous arrival. But on the present account, that description has things backwards. Between jump and arrival you are *not* accelerating. You are in a state of free fall and moving (approximately) along a spacetime geodesic. But before the jump, and after the arrival, you *are* accelerating. The floor of the observation desk, and then later the sidewalk, push you away from a geodesic path. The all-important idea here is that we are incorporating the “gravitational field” into the geometric structure of spacetime, and particles traverse geodesics if and only if they are acted upon by no forces “except gravity”.

The acceleration of any massive particle, i.e. its deviation from a geodesic trajectory, is determined by the forces acting on it (other than “gravity”). If the particle has mass $m > 0$, and the vector field F^a on $\gamma[I]$ represents the vector sum of the various (non-gravitational) forces acting on the particle, then the particle’s four-acceleration $\xi^n \nabla_n \xi^a$ satisfies:

$$(15) \quad F^a = m \xi^n \nabla_n \xi^a.$$

This is our version of Newton’s second law of motion.

Consider an example. Electromagnetic fields are represented by smooth, anti-symmetric fields F_{ab} . (Here “anti-symmetry” is the condition that $F_{ba} = -F_{ab}$.) If a particle with mass $m > 0$, charge q , and four-velocity field ξ^a is present, the force exerted by the field on the particle at a point is given by $q F^a_b \xi^b$. If we use this expression for the left side of (15), we arrive at the Lorentz law of motion for charged particles in the presence of an electromagnetic field:

$$(16) \quad q F^a_b \xi^b = m \xi^b \nabla_b \xi^a.^{20}$$

2.4 Matter Fields

In classical relativity theory, one generally takes for granted that all that there is, and all that happens, can be described in terms of various *matter fields*, e.g. material fluids and electromagnetic fields.²¹ Each such field is represented by one or more smooth tensor (or spinor) fields on the spacetime manifold M . Each is assumed to satisfy field equations involving the fields that represent it and the spacetime metric g_{ab} .

For present purposes, the most important basic assumption about the matter fields is the following.

Associated with each matter field \mathcal{F} is a symmetric smooth tensor field T_{ab} characterized by the property that, for all points p in M , and all future-directed, unit timelike vectors ξ^a at p , $T^a_b \xi^b$ is the four-momentum density of \mathcal{F} at p as determined relative to ξ^a .

T_{ab} is called the *energy-momentum* field associated with \mathcal{F} . The four-momentum density vector $T^a_b \xi^b$ at p can be further decomposed into its temporal and spatial components relative to ξ^a , just as the four-momentum of a massive particle was decomposed in the preceding section. The coefficient of ξ^a in the first component, $T_{ab} \xi^a \xi^b$, is the *energy density* of \mathcal{F} at p as determined relative to ξ^a . The second component, $T_{nb} (g^{an} - \xi^a \xi^n) \xi^b$, is the *three-momentum density* of \mathcal{F} at p as determined relative to ξ^a .

Other assumptions about matter fields can be captured as constraints on the energy-momentum tensor fields with which they are associated. Examples are the following. (Suppose T_{ab} is associated with matter field \mathcal{F} .)

Weak Energy Condition: Given any future-directed unit timelike vector ξ^a at any point in M , $T_{ab} \xi^a \xi^b \geq 0$.

Dominant Energy Condition: Given any future-directed unit timelike vector ξ^a at any point in M , $T_{ab} \xi^a \xi^b \geq 0$ and $T^a_b \xi^b$ is timelike or null.

Conservation Condition: $\nabla_a T^{ab} = 0$ at all points in M .

The first asserts that the energy density of \mathcal{F} , as determined by any observer at any point, is non-negative. The second adds the requirement that the four-momentum density of \mathcal{F} , as determined by any observer at any point, is a future-directed causal (i.e. timelike

²¹This being the case, the question arises how (or whether) one can adequately recover talk about “point particles” in terms of the matter fields. We will say just a bit about the question in this section.

or null) vector. The addition can be understood as the assertion that there is an upper bound to the speed with which energy-momentum can propagate (as determined by any observer). It captures something of the flavor of principle C1 in section 2.1, but avoids reference to “point particles”.²²

The conservation condition, finally, asserts that the energy-momentum carried by \mathcal{F} is locally conserved. If two or more matter fields are present in the same region of spacetime, it need not be the case that each one individually satisfies the condition. Interaction may occur. But it *is* a fundamental assumption that the composite energy-momentum field formed by taking the sum of the individual ones satisfies it. Energy-momentum can be transferred from one matter field to another, but it cannot be created or destroyed.

The dominant energy and conservation conditions have a number of joint consequences that support the interpretations just given. We mention two. The first requires a preliminary definition.

Let (M, g_{ab}) be a fixed relativistic spacetime, and let S be an achronal subset of M (i.e. a subset in which there do *not* exist points p and q such that $p \ll q$). The *domain of dependence* $D(S)$ of S is the set of all points p in M with this property: given any smooth causal curve without (past or future) endpoint,²³ if (its image) passes through p , then it necessarily intersects S . For all standard matter fields, at least, one can prove a theorem to the effect that “what happens on S fully determines what happens throughout $D(S)$ ”. (See Earman (this volume, chapter 15).) Here we consider just a special case.

PROPOSITION 6. *Let S be an achronal subset of M . Further let T_{ab} be a smooth symmetric field on M that satisfies both the dominant energy and conservation conditions. Finally, assume $T_{ab} = \mathbf{0}$ on S . Then $T_{ab} = \mathbf{0}$ on all of $D(S)$.*

The intended interpretation of the proposition is clear. If energy-momentum cannot propagate (locally) outside the null-cone, and if it is conserved, and if it vanishes on S , then it must vanish throughout $D(S)$. After all, how could it “get to” any point in $D(S)$? Note that our formulation of the proposition does not presuppose any particular physical interpretation of the symmetric field T_{ab} . All that is required is that it satisfy the two stated conditions. (For a proof, see Hawking and Ellis [1972, p. 94].)

The next proposition (Geroch and Jang [1975]) shows that, in a sense, if one assumes the dominant energy condition and the conservation condition, then one can *prove* that free massive point particles traverse the images of timelike geodesics. (Recall principle P1 in section 2.3.) The trick is to find a way to talk about “point particles” in the language of extended matter fields.

PROPOSITION 7. *Let $\gamma: I \rightarrow M$ be smooth curve. Suppose that given any open subset O of M containing $\gamma[I]$, there exists a smooth symmetric field T_{ab} on M such that:*

- (1) T_{ab} satisfies the dominant energy condition;
- (2) T_{ab} satisfies the conservation condition;

²²This is the standard formulation of the dominant energy condition. The fit with C1 would be even closer if we strengthened the condition slightly so as to be appropriate, specifically, for massive matter fields: at any point p in M , if $T^a_b \neq \mathbf{0}$ there, then $T^a_b \xi^b$ is *timelike* for all future-directed unit timelike vectors ξ^a at p .

²³Let $\gamma: I \rightarrow M$ be a smooth curve. We say that a point p in M is a *future-endpoint* of γ if, for all open sets O containing p , there exists an s_0 in I such that for all $s \in I$, if $s \geq s_0$, then $\gamma(s) \in O$, i.e. the image of γ eventually enters and remains in O . (*Past-endpoints* are defined similarly.)

(3) $T_{ab} = \mathbf{0}$ outside of O ;

(4) $T_{ab} \neq \mathbf{0}$ at some point in O .

Then γ is timelike, and can be reparametrized so as to be a geodesic.

The proposition might be paraphrased this way. If a smooth curve in spacetime is such that arbitrarily small free bodies could contain the image of the curve in their worldtubes, then the curve must be a timelike geodesic (up to reparametrization). In effect, we are trading in “point particles” in favor of nested convergent sequences of smaller and smaller extended particles. (Bodies here are understood to be “free” if their internal energy-momentum is conserved. If a body is acted upon by a field, it is only the composite energy-momentum of the body and field together that is conserved.)

Note that our formulation of the proposition takes for granted that we can keep the background spacetime structure (M, g_{ab}) fixed while altering the fields T_{ab} that live on M . This is justifiable only to the extent that, in each case, T_{ab} is understood to represent a test body whose effect on the background spacetime structure is negligible.²⁴ Note also that we do not have to assume at the outset that the curve γ is timelike. That follows from the other assumptions.

We have here a precise proposition in the language of matter fields that, at least to some degree, captures principle P1 (concerning the behavior of free massive point particles). Similarly, it is possible to capture C2 (concerning the behavior of light) with a proposition about the behavior of solutions to Maxwell’s equations in a limiting regime (“the geometrical limit”) where wavelengths are small. It asserts, in effect, that when one passes to this limit, packets of electromagnetic waves are constrained to move along (images of) null geodesics. (See Wald [1984, p. 71].)

Now we consider an example. *Perfect fluids* are represented by three objects: a four-velocity field η^a , an energy density field ρ , and an isotropic pressure field p (the latter two as determined by a “co-moving” observer at rest in the fluid). In the special case where the pressure p vanishes, one speaks of a *dust field*. Particular instances of perfect fluids are characterized by “equations of state” that specify p as a function of ρ . (Specifically excluded here are such complicating factors as anisotropic pressure, shear stress, and viscosity.) Though ρ is generally assumed to be non-negative (see below), some perfect fluids (e.g. to a good approximation, water) can exert negative pressure. The energy-momentum tensor field associated with a perfect fluid is:

$$(17) \quad T_{ab} = \rho \eta_a \eta_b - p (g_{ab} - \eta_a \eta_b).$$

Notice that the energy-momentum density vector of the fluid at any point, as determined by a co-moving observer (i.e. as determined relative to η^a), is $T^a_b \eta^b = \rho \eta^a$. So we can understand ρ , equivalently, as the energy density of the fluid relative to η^a , i.e. $T_{ab} \eta^a \eta^b$, or as the (rest) mass density of the fluid, i.e. the length of $\rho \eta^a$. (Of course, the situation here corresponds to that of a point particle with mass m and four-velocity η^a , as considered in section 2.3.)

²⁴Stronger theorems have been proved (see Ehlers and Geroch [2004]) in which it is not required that the perturbative effect of the extended body disappear entirely at each stage of the limiting process, but only that, in a certain sense, it disappear in the limit.

In the case of a perfect fluid, the weak energy condition (WEC), dominant energy condition (DEC), and conservation condition (CC) come out as follows.

$$\text{WEC} \iff \rho \geq 0 \quad \text{and} \quad p \geq -\rho$$

$$\text{DEC} \iff \rho \geq 0 \quad \text{and} \quad \rho \geq p \geq -\rho$$

$$\text{CC} \iff \begin{cases} (\rho + p) \eta^b \nabla_b \eta^a - (g^{ab} - \eta^a \eta^b) \nabla_b p & = \mathbf{0} \\ \eta^b \nabla_b \rho + (\rho + p) (\nabla_b \eta^b) & = 0. \end{cases}$$

Consider the two equations jointly equivalent to the conservation condition. The first is the equation of motion for a perfect fluid. We can think of it as a relativistic version of Euler’s equation. The second is an equation of continuity (or conservation) in the sense familiar from classical fluid mechanics. It is easiest to think about the special case of a dust field ($p = 0$). In that case, the equation of motion reduces to the geodesic equation: $\eta^b \nabla_b \eta^a = \mathbf{0}$. That makes sense. In the absence of pressure, particles in the fluid are free particles. And the conservation equation reduces to: $\eta^b \nabla_b \rho + \rho (\nabla_b \eta^b) = 0$. The first term gives the instantaneous rate of change of the fluid’s energy density, as determined by a co-moving observer. The term $\nabla_b \eta^b$ gives the instantaneous rate of change of its volume, per unit volume, as determined by that observer. In a more familiar notation, the equation might be written $\frac{d\rho}{ds} + \frac{\rho}{V} \frac{dV}{ds} = 0$ or, equivalently, $\frac{d(\rho V)}{ds} = 0$. (Here we use s for elapsed proper time.) It asserts that (in the absence of pressure, as determined by a co-moving observer) the energy contained in an (infinitesimal) fluid blob remains constant, even as its volume changes.

In the general case, the situation is more complex because the pressure in the fluid contributes to its energy (as determined relative to particular observers), and hence to what might be called its “effective mass density”. (If you compress a fluid blob, it gets heavier.) In this case, the WEC comes out as the requirement that $(\rho + p) \geq 0$ in addition to $\rho \geq 0$. If we take $h^{ab} = (g^{ab} - \eta^a \eta^b)$, the equation of motion can be expressed as:

$$(\rho + p) \eta^b \nabla_b \eta^a = h^{ab} \nabla_b p.$$

This is an instance of the “second law of motion” (15) as applied to an (infinitesimal) blob of fluid. On the left we have: “effective mass density \times acceleration”. On the right, we have the force acting on the blob. We can think of it as minus²⁵ the gradient of the pressure (as determined by a co-moving observer). Again, this makes sense. If the pressure on the left side of the blob is greater than that on the right, it will move to the right. The presence of the non-vanishing term $(p \nabla_b \eta^b)$ in the conservation equation is now required because the energy of the blob is *not* constant when its volume changes as a result of the pressure. The equation governs the contribution made to its energy by pressure.

2.5 Einstein’s Equation

Once again, let (M, g_{ab}) be our background relativistic spacetime with a specified temporal orientation.

²⁵The minus sign comes in because of our sign conventions.

It is one of the fundamental ideas of relativity theory that spacetime structure is not a fixed backdrop against which the processes of physics unfold, but instead participates in that unfolding. It posits a dynamical interaction between the spacetime metric in any region and the matter fields there. The interaction is governed by *Einstein's field equation*

$$(18) \quad R_{ab} - \frac{1}{2} R g_{ab} - \lambda g_{ab} = 8\pi T_{ab},$$

or, equivalently,

$$(19) \quad R_{ab} = 8\pi (T_{ab} - \frac{1}{2} T g_{ab}) - \lambda g_{ab}.$$

Here λ is the *cosmological constant*, R_{ab} ($= R^n_{abn}$) is the *Ricci tensor field*, R ($= R^a_a$) is the Riemann scalar curvature field, and T is the contracted field T^a_a .²⁶ We start with four remarks about (18), and then consider an alternative formulation that provides a geometric interpretation of sorts.

(1) It is sometimes taken to be a version of "Mach's principle" that "the spacetime metric is uniquely determined by the distribution of matter". And it is sometimes proposed that the principle can be captured in the requirement that "if one first specifies the energy-momentum distribution T_{ab} on the spacetime manifold M , then there is exactly one (or at most one) Lorentzian metric g_{ab} on M that, together with T_{ab} , satisfies (18)". But there is a serious problem with the proposal. In general, one *cannot* specify the energy-momentum distribution in the absence of a spacetime metric. (E.g. one cannot have a notion of energy-density unless one has a notion of volume.) Indeed, in typical cases the metric enters explicitly in the expression for T_{ab} . (Recall the expression (17) for a perfect fluid.) Thus, in looking for solutions to (18), one must, in general, solve simultaneously for the metric *and* matter field distribution.

(2) Given any smooth metric g_{ab} on M , there certainly exists a smooth symmetric field T_{ab} on M that, together with g_{ab} , is a solution to (18). It suffices to define T_{ab} by the left side of the equation. But the field T_{ab} so introduced will not, in general, be the energy-momentum field associated with any known matter field. It will not even satisfy the weak energy condition discussed in section 2.4. With the latter constraint on T_{ab} in place, Einstein's equation is an entirely non-trivial restriction on spacetime structure.

Discussions of spacetime structure in classical relativity theory proceed on three levels according to the stringency of the constraints imposed on T_{ab} . At the first level, one considers only "exact solutions", i.e. solutions where T_{ab} is, in fact, the aggregate energy-momentum field associated with one or more known matter fields. So, for example, one might undertake to find all perfect fluid solutions exhibiting particular symmetries. At the second level, one considers the larger class of what might be called "generic solutions", i.e. solutions where T_{ab} satisfies one or more generic constraints (of which the weak and dominant energy conditions are examples). It is at this level, for example, that the singularity theorems of Penrose and Hawking (Hawking and Ellis [1972]) are proved. Finally, at the third level, one drops all restrictions on T_{ab} , and Einstein's equation plays no role. Many results about global structure are proved at this level, e.g. the assertion

²⁶We use "geometrical units" in which the gravitational constant G , as well as the speed of light c , is 1.

that there exist closed timelike curves in any relativistic spacetime (M, g_{ab}) where M is compact.

(3) The role played by the cosmological constant in Einstein’s equation remains a matter of controversy. Einstein initially added the term $(-\lambda g_{ab})$ in 1917 to allow for the possibility of a static cosmological model (which, at the time, was believed necessary to properly represent the actual universe).²⁷ But there were clear problems with doing so. In particular, one does not recover Poisson’s equation (the field equation of Newtonian gravitation theory) as a limiting form of Einstein’s equation unless $\lambda = 0$. (See point (4) below.) Einstein was quick to revert to the original form of the equation after Hubble’s redshift observations gave convincing evidence that the universe is, in fact, expanding. (That the theory suggested the possibility of cosmic expansion before those observations must count as one of its great successes.) Since then the constant has often been reintroduced to help resolve discrepancies between theoretical prediction and observation, and then abandoned when the (apparent) discrepancies were resolved. The controversy continues. Recent observations indicating an accelerating rate of cosmic expansion have led many cosmologists to believe that our universe is characterized by a positive value for λ . (See Earman [2001] for an overview.)

Claims about the value of the cosmological constant are sometimes cast as claims about the “energy-momentum content of the vacuum”. This involves bringing the term $(-\lambda g_{ab})$ from the left side of equation (18) to the right, and re-interpreting it as an energy-momentum field, i.e. taking Einstein’s equation in the form

$$(20) \quad R_{ab} - \frac{1}{2} R g_{ab} = 8\pi (T_{ab} + T_{ab}^{VAC}),$$

where $T_{ab}^{VAC} = \frac{\lambda}{8\pi} g_{ab}$. Here T_{ab} is still understood to represent the aggregate energy-momentum of all normal matter fields. But T_{ab}^{VAC} is now understood to represent the residual energy-momentum associated with empty space. Given any unit timelike vector ξ^a at a point, $(T_{ab}^{VAC} \xi^a \xi^b)$ is $\frac{\lambda}{8\pi}$. So, on this re-interpretation, λ comes out (up to the factor 8π) as the energy-density of the vacuum as determined by *any* observer, at any point in spacetime.

It should be noted that there is a certain ambiguity involved in referring to λ as the cosmological *constant* (and a corresponding ambiguity as to what counts as a solution to Einstein’s equation). We can take (M, g_{ab}, T_{ab}) to qualify if it satisfies the equation for some value (or other) of λ . Or, more stringently, one can take it to qualify if it satisfies the equation for some value of λ that is fixed, once and for all, i.e. the same for all models (M, g_{ab}, T_{ab}) . In effect, we have here two versions of “relativity theory”. (See Earman [2003] for discussion of what is at stake in choosing between the two.)

(4) It is instructive to consider the relation of Einstein’s equation to Poisson’s equation, the field equation of Newtonian gravitation theory:

$$(21) \quad \nabla^2 \phi = 4\pi \rho.$$

²⁷He did so for other reasons as well (see Earman [2001]), but I will pass over them here.

Here ϕ is the Newtonian gravitational potential, and ρ is the Newtonian mass density function. In the “geometrized” formulation of the theory that we will consider in section 3.2, one trades in the potential ϕ in favor of a curved derivative operator, and Poisson’s equation comes out as

$$(22) \quad R_{ab} = 4\pi\rho t_{ab},$$

where R_{ab} is the Ricci tensor field associated with the new curved derivative operator, and t_{ab} is the temporal metric.

The geometrized formulation of Newtonian gravitation was discovered after general relativity (in the 1920s). But now, after the fact, we can put ourselves in the position of a hypothetical investigator who is considering possible candidates for a relativistic field equation, and knows about the geometrized formulation of Newtonian theory. What could be more natural than the attempt to adopt or adapt (22)? In the empty space case ($\rho = 0$), this strategy suggests the equation $R_{ab} = \mathbf{0}$, which is, of course, Einstein’s equation (19) for $T_{ab} = \mathbf{0}$ and $\lambda = 0$. This seems to me, by far, the best route to the latter equation. Start with the Newtonian empty space equation ($R_{ab} = \mathbf{0}$) and then simply leave it intact!

No such simple extrapolation is possible in the general case ($\rho \neq 0$). Indeed, I know of no heuristic argument for the full version of Einstein’s equation (with or without cosmological constant) that is nearly so convincing. But one can try something like the following. The closest counterparts to (22) would seem to be ones of the form: $R_{ab} = 4\pi K_{ab}$, where K_{ab} is a symmetric tensorial function of T_{ab} and g_{ab} . The possibilities for K_{ab} include T_{ab} , $g_{ab}T$, $T_a^m T_{mb}$, $g_{ab}(T^{mn}T_{mn})$, ..., and linear combinations of these terms. All but the first two involve terms that are second order or higher in T_{ab} . So, for example, in the special case of a dust field with energy density ρ and four-velocity η^a , they will contain occurrences of ρ^n with $n \geq 2$. (E.g. $g_{ab}(T^{mn}T_{mn})$ comes out as $\rho^2 g_{ab}$.) But, presumably, only terms first order in ρ should appear if the equation is to have a proper Newtonian limit. This suggests that we look for a field equation of the form

$$(23) \quad R_{ab} = 4\pi [k T_{ab} + l g_{ab}T]$$

or, equivalently,²⁸

$$(24) \quad R_{ab} - \frac{l}{(k + 4l)} R g_{ab} = 4\pi k T_{ab},$$

for some real numbers k and l . Let $G_{ab}(k, l)$ be the field on the left side of the equation. It follows from the conservation condition that the field on the right side is divergence free, i.e. $\nabla_a (4\pi k T^{ab}) = \mathbf{0}$. So the conservation condition and (24) can hold jointly only if

$$\nabla_a G^{ab}(k, l) = \mathbf{0}.$$

But by the “Bianchi identity”(Wald [1984, pp. 39-40]),

$$(25) \quad \nabla_a (R^{ab} - \frac{1}{2} R g^{ab}) = \mathbf{0}.$$

²⁸Contraction on ‘a’ and ‘b’ in (23) yields: $R = 4\pi(k + 4l)T$. Solving for T , and substituting for T in (23) yields (24).

The latter two conditions imply

$$\left[\frac{l}{(k + 4l)} - \frac{1}{2} \right] \nabla_a (Rg^{ab}) = \mathbf{0}.$$

Now $\nabla_a (Rg^{ab}) = \mathbf{0}$ is an unreasonable constraint.²⁹ So the initial scalar term must be 0. Thus, we are left with the conclusion that the conservation condition and (24) can hold jointly only if $k + 2l = 0$, in which case (23) reduces to

$$(26) \quad R_{ab} = 4\pi k \left[T_{ab} - \frac{1}{2} g_{ab} T \right].$$

It remains to argue that k must be 2 if (26) is to have a proper Newtonian limit. To do so, we consider, once again, the special case of a dust field with energy density ρ and four-velocity η^a . Then, $T_{ab} = \rho \eta_a \eta_b$, and $T = \rho$. If we insert these values in (26) and contract with $\eta^a \eta^b$, we arrive at

$$(27) \quad R_{ab} \eta^a \eta^b = 2\pi k \rho.$$

Now the counterpart to a four-velocity field in Newtonian theory is a vector field of unit temporal length, i.e. a field η^a where $t_{ab} \eta^a \eta^b = 1$. If we contract the geometrized version of Poisson's equation (22) with $\eta^a \eta^b$, we arrive at: $R_{ab} \eta^a \eta^b = 4\pi \rho$. Comparing this expression for $R_{ab} \eta^a \eta^b$ with that in (27), we are led to the conclusion that $k = 2$, in which case (26) is just Einstein's equation (19) with $\lambda = 0$.

Summarizing now, we have suggested that if one starts with the geometrized version of Poisson's equation (22) and looks for a relativistic counterpart, one is plausibly led to Einstein's equation with $\lambda = 0$. It is worth noting that if we had started instead with a variant of (22) incorporating a "Newtonian cosmological constant"

$$(28) \quad R_{ab} + \lambda t_{ab} = 4\pi \rho t_{ab},$$

we would have been led instead to Einstein's equation (19) without restriction on λ . We can think of (28) as the geometrized version of

$$(29) \quad \nabla^2 \phi + \lambda = 4\pi \rho.$$

Let's now put aside the question of how one might try to motivate Einstein's equation. However one arrives at it, the equation — let's now take it in the form (18) — can be understood to assert a dynamical connection between a certain tensorial measure of spacetime curvature (on the left side) and the energy-momentum tensor field (on the right side). It turns out that one can reformulate the connection in a way that makes reference only to scalar quantities, as determined relative to arbitrary observers. The reformulation provides a certain insight into the geometric significance of the equation.³⁰

²⁹It implies that R is constant and, hence, if (23) holds, that T is constant (since (23) implies $R = 4\pi(k + 4l)T$). But this, in turn, is an unreasonable constraint on the energy-momentum distribution T_{ab} . E.g. in the case of a dust field with $T_{ab} = \rho \eta^a \eta^b$, $T = \rho$, and so the constraint implies that ρ is constant. This is unreasonable since it rules out any possibility of cosmic expansion. (Recall the discussion toward the close of section 2.4.)

³⁰Another approach to its geometrical significance proceeds via the equation of geodesic deviation. See, for example, Sachs and Wu [1977b, p. 114].

Let S be any smooth spacelike hypersurface in M .³¹ The background metric g_{ab} induces a (three-dimensional) metric ${}^3g_{ab}$ on S . In turn, this metric determines on S a derivative operator, an associated Riemann curvature tensor field ${}^3R^a_{bcd}$, and a scalar curvature field ${}^3R = ({}^3R^a_{bca})({}^3g^{bc})$. Our reformulation of Einstein’s equation will direct attention to the values of 3R at a point for a particular family of spacelike hypersurfaces passing through it.³²

Let p be any point in M and let ξ^a be any future-directed unit timelike vector at p . Consider the set of all geodesics through p that are orthogonal to ξ^a there. The (images of these) curves, at least when restricted to a sufficiently small open set containing p , sweep out a smooth spacelike hypersurface S .³³ (See figure 2.) We will call it a *geodesic*




Figure 2. A “geodesic hypersurface” through a point is constructed by projecting geodesics in all directions orthogonal to a given timelike vector there.

hypersurface. (We cannot speak of *the* geodesic hypersurface through p orthogonal to ξ^a because we have left open how far the generating geodesics are extended. But given any two, their restrictions to a suitably small open set containing p coincide.)

Geodesic hypersurfaces are of interest in their own right, the present context aside, because they are natural candidates for a notion of “local simultaneity slice” (relative to a timelike vector at a point). What matters here, though, is that, by the first Gauss-Codazzi equation (Wald [1984, p. 258]), we have

$$(30) \quad {}^3R = R - 2 R_{ab}\xi^a\xi^b$$

at p .³⁴ Here we have expressed the (three-dimensional) Riemann scalar curvature of S at

³¹We can take this to mean that S is a smooth, imbedded, three-dimensional submanifold of M with the property that any curve $\gamma: I \rightarrow M$ with image in S is spacelike.

³²In the case of a surface in three-dimensional Euclidean space, the associated Riemann scalar curvature 2R is (up to a constant) just ordinary Gaussian surface curvature. We can think of 3R in the present context as a higher dimensional analogue that gives averaged values of Gaussian surface curvature. This can be made precise. See, for example, Laugwitz [1965, p. 127].

³³More precisely, let S_p be the spacelike hyperplane in M_p orthogonal to ξ^a . Then for any sufficiently small open set O in M_p containing p , the image of $(S_p \cap O)$ under the exponential map $exp: O \rightarrow M$ is a smooth spacelike hypersurface. We can take it to be S . (See, for example, Hawking and Ellis [1972, p. 33].)

³⁴Let ξ^a — we use the same notation — be the extension of the original vector at p to a smooth future-directed unit timelike vector field on S that is everywhere orthogonal to S . Then the first Gauss-Codazzi equation asserts that at all points of S

$${}^3R = R - 2 R_{ab}\xi^a\xi^b + \pi_{ab} h^{ab} + \pi_{ab} \pi^{ab},$$

where h_{ab} is the spatial projection field $(g_{ab} - \xi_a\xi_b)$ on S , and π_{ab} is the *extrinsic curvature* field $\frac{1}{2} \mathcal{L}_\xi h_{ab}$ on S . But our construction guarantees that π_{ab} vanish at p .

p in terms of the (four-dimensional) Riemann scalar curvature of M at p and the Ricci tensor there. And so, if Einstein's equation (18) holds, we have

$$(31) \quad {}^3R = -16\pi (T_{ab} \xi^a \xi^b) - 2\lambda.$$

at p .

One can also easily work backwards to recover Einstein's equation at p from the assumption that (31) holds for *all* unit timelike vectors ξ^a at p (and all geodesic hypersurfaces through p orthogonal to ξ^a). Thus, we have the following equivalence.

PROPOSITION 8. *Let T_{ab} be a smooth symmetric field on M , and let p be a point in M . Then Einstein's equation $R_{ab} - \frac{1}{2} R g_{ab} - \lambda g_{ab} = 8\pi T_{ab}$ holds at p iff for all future-directed unit timelike vectors ξ^a at p , and all geodesic hypersurfaces through p orthogonal to ξ^a , the scalar curvature 3R of S satisfies ${}^3R = [-16\pi (T_{ab} \xi^a \xi^b) - 2\lambda]$ at p .*

The result is particularly instructive in the case where $\lambda = 0$. Then (31) directly equates an intuitive scalar measure of spatial curvature (as determined relative to ξ^a) with energy density (as determined relative to ξ^a).

2.6 Congruences of Timelike Curves and "Public Space"

In this section, we consider congruences of timelike curves. (We understand these to be sets of timelike curves that "fill" a region of spacetime in the sense that exactly one curve (image) in the set passes through each point in the region.) We think of them as representing the worldlines of a dense swarm of particles or the elements of a fluid.

Each such congruence is generated by a future-directed, unit timelike vector field (that represents the four-velocity field of our particle swarm or fluid). We work directly with these generating fields in what follows.

Once again, let (M, g_{ab}) be our background relativistic spacetime (endowed with a temporal orientation). Let ξ^a be a smooth, future-directed, unit timelike vector field on M (or some open subset thereof). Finally, h_{ab} be the spatial projection field determined by ξ^a .

The *rotation* and *expansion* tensor fields associated with ξ^a are defined as follows:

$$(32) \quad \omega_{ab} = h_{[a}{}^m h_{b]}{}^n \nabla_m \xi_n$$

$$(33) \quad \theta_{ab} = h_{(a}{}^m h_{b)}{}^n \nabla_m \xi_n.$$

They are smooth fields, orthogonal to ξ^a in both indices, and satisfy

$$(34) \quad \nabla_a \xi_b = \omega_{ab} + \theta_{ab} + \xi_a (\xi^n \nabla_n \xi_b).$$

We can give the two fields ω_{ab} and θ_{ab} a geometric interpretation. Let η^a be a vector field on the worldline of a particle O that is "carried along by the flow of ξ^a ", i.e. $\mathcal{L}_\xi \eta^a = 0$, and is orthogonal to ξ^a at a point p . (Here $\mathcal{L}_\xi \eta^a$ is the Lie derivative of η^a with respect to ξ^a .³⁵) We can think of η^a at p as a spatial "connecting vector" that spans the distance

³⁵We drop the index on ξ here to avoid giving the impression that $\mathcal{L}_\xi g_{ab}$ is a three index tensor field. Lie derivatives are always taken with respect to (contravariant) vector fields, so no ambiguity is introduced when the index is dropped.

between O and a neighboring particle N that is “infinitesimally close”. The instantaneous velocity of N relative to O at p is given by $\xi^n \nabla_n \eta^a$. But $\xi^n \nabla_n \eta^a = \eta^n \nabla_n \xi^a$ (since $\mathcal{L}_\xi \eta^a = 0$). So, by (34), and the orthogonality of ξ^a with η^a at p , we have

$$(35) \quad \xi^n \nabla_n \eta^a = (\omega_n^a + \theta_n^a) \eta^n.$$

at the point. Here we have simply decomposed the relative velocity vector into two components. The first, $(\omega_n^a \eta^n)$, is orthogonal to η^a (since ω_{ab} is anti-symmetric). It gives the instantaneous *rotational velocity* of N with respect to O at p .

In support of this interpretation, consider the instantaneous rate of change of the squared length $(-\eta^a \eta_a)$ of η^a at p . It follows from (35) that

$$(36) \quad \xi^n \nabla_n (-\eta^a \eta_a) = -2 \theta_{na} \eta^n \eta^a.$$

Thus the computed rate of change depends solely on θ_{ab} . Suppose $\theta_{ab} = 0$. Then the instantaneous velocity of N with respect to O at p has vanishing radial component. If $\omega_{ab} \neq 0$, N still exhibits a non-zero velocity with respect to O . But it can only be a rotational velocity. The two conditions ($\theta_{ab} = 0$ and $\omega_{ab} \neq 0$) jointly characterize “rigid rotation”.

The condition $\omega_{ab} = 0$, by itself, characterizes *irrotational* flow. One gains considerable insight into the condition by considering a second, equivalent formulation. Let us say that the field ξ^a is *hypersurface orthogonal* if there exist smooth, real valued maps f and g (with the same domains of definition as ξ^a) such that, at all points, $\xi_a = f \nabla_a g$. Note that if the condition is satisfied, then the hypersurfaces of constant g value are everywhere orthogonal to ξ^a .³⁶ Let us further say that ξ^a is *locally hypersurface orthogonal* if the restriction of ξ^a to every sufficiently small open set is hypersurface orthogonal.

PROPOSITION 9. *Let ξ^a be a smooth, future-directed unit timelike vector field defined on M (or some open subset of M). Then the following conditions are equivalent.*

- (1) $\omega_{ab} = 0$ everywhere.
- (2) ξ^a is locally hypersurface orthogonal.

The implication from (2) to (1) is immediate.³⁷ But the converse is non-trivial. It is a special case of Frobenius’s theorem (Wald [1984, p. 436]). The qualification ‘locally’ can be dropped in (2) if the domain of ξ^a is, for example, simply connected.

There is a nice picture that goes with the proposition. Think about an ordinary rope. In its natural twisted state, the rope cannot be sliced by an infinite family of slices in such a way that each slice is orthogonal to all fibers. But if the rope is first untwisted, such a

³⁶For if η^a is a vector tangent to one of these hypersurfaces, $\eta^n \nabla_n g = 0$. So $\eta^n \xi_n = \eta^n (f \nabla_n g) = 0$.

³⁷Assume that $\xi_a = f \nabla_a g$. Then

$$\begin{aligned} \omega_{ab} &= h_{[a}^m h_{b]}^n \nabla_m \xi_n = h_{[a}^m h_{b]}^n \nabla_m (f \nabla_n g) \\ &= f h_{[a}^m h_{b]}^n \nabla_m \nabla_n g + h_{[a}^m h_{b]}^n (\nabla_m f) (\nabla_n g) \\ &= f h_a^m h_b^n \nabla_{[m} \nabla_{n]} g + h_a^m h_b^n (\nabla_{[m} f) (\nabla_{n]} g). \end{aligned}$$

But $\nabla_{[m} \nabla_{n]} g = 0$ since ∇_a is torsion-free, and the second term in the final line vanishes as well since $h_b^n \nabla_n g = f^{-1} h_b^n \xi_n = 0$. So $\omega_{ab} = 0$.

slicing is possible. Thus orthogonal sliceability is equivalent to fiber untwistedness. The proposition extends this intuitive equivalence to the four-dimensional “spacetime ropes” (i.e. congruences of worldlines) encountered in relativity theory. It asserts that a congruence is irrotational (i.e. exhibits no twistedness) iff it is, at least locally, hypersurface orthogonal

Suppose that our vector field ξ^a is irrotational and, to keep things simple, suppose that its domain of definition is simply connected. Then the hypersurfaces to which it is orthogonal are natural candidates for constituting “space” at a given “time” relative to ξ^a or, equivalently, relative to its associated set of integral curves. This is a notion of *public space* to be contrasted with *private space*, which is determined relative to individual timelike vectors or timelike curves.³⁸ Perhaps the best candidates for the latter are the “geodesic hypersurfaces” we considered, in passing, in section 2.5. (Given a point p and a timelike vector ξ^a there, we took a “geodesic hypersurface through p orthogonal to ξ^a ” to be a spacelike hypersurface generated by geodesics through p orthogonal to ξ^a .)

The distinction between public and private space is illustrated in Figure 3. There we




Figure 3. “Private space” $S_{private}$ at p relative to L , and “public space” S_{public} at p relative to a congruence of timelike curves of which L is a member.

consider a congruence of future-directed timelike half-geodesics in Minkowski spacetime starting at some particular point p . One line L in the congruence is picked out along with a point q on it. Private space relative to L at q is a spacelike hypersurface $S_{private}$ that is flat, i.e. the metric induced on $S_{private}$ has a Riemann curvature tensor field ${}^3R^a_{bcd}$ that vanishes everywhere. In contrast, public space at q relative to the congruence is a spacelike hypersurface S_{public} of constant negative curvature. If ξ^a is the future-directed unit timelike vector field everywhere tangent to the congruence, and $h_{ab} = (g_{ab} - \xi_a \xi_b)$ is its associated spatial projection field, then the curvature tensor field on S_{public} associated with h_{ab} has the form ${}^3R_{abcd} = -\frac{1}{K^2}(h_{ac} h_{bd} - h_{ad} h_{bc})$, where K is the distance along L from p to q . (This is the characteristic form for a three-manifold of constant curvature $-\frac{1}{K^2}$.)

We have been considering “public space” as determined relative to an irrotational congruence of timelike curves. There is another sense in which one might want to use the term. Consider, for example, “geometry on the surface of a rigidly rotating disk” in Minkowski spacetime. (There is good evidence that Einstein’s realization that this ge-

³⁸The distinction between “public space” and “private space” is discussed in Rindler [1981] and Page [1983]. The terminology is due to E. A. Milne.

ometry is non-Euclidean played an important role in his development of relativity theory (Stachel [1980]).) One needs to ask in what sense the surface of a rotating disk *has* a geometric structure.

We can certainly model the rigidly rotating disk as a congruence of timelike curves in Minkowski spacetime. (Since the disk is two-dimensional, the congruence will be confined to a three-dimensional, timelike submanifold M' of M .) But precisely because the disk is rotating, we cannot find hypersurfaces everywhere orthogonal to the curves and understand the geometry of the disk to be the geometry induced on them — or, strictly speaking, induced on the two-dimensional manifolds determined by the intersection of the putative hypersurfaces with M' — by the background spacetime metric g_{ab} .

The alternative is to think of “space” as constituted by the “manifold of trajectories”, i.e. take the individual timelike curves in the congruence to play the role of spatial points, and consider the metric induced on *this* manifold by the background spacetime metric. The construction will not work for an arbitrary congruence of timelike curves. It is essential that we are dealing here with a “stationary” system. (The metric induced on the manifold of trajectories (when the construction works) is fixed and frozen.) But it *does* work for these systems, at least. More precisely, anticipating the terminology of the following section, it works if the four-velocity field of the congruence in question is proportional to a Killing field. (The construction is presented in detail in Geroch [1971, Appendix].)

Thus we have two notions of “public space”. One is available if the four-velocity field of the congruence in question is irrotational; the other if it is proportional to a Killing field. Furthermore, if the four-velocity field is irrotational *and* proportional to Killing field, as is the case when we dealing with a “static” system, then the two notions of public space are essentially equivalent.

2.7 Killing Fields and Conserved Quantities

Let κ^a be a smooth vector field on M . We say it is a *Killing field* if $\mathcal{L}_\kappa g_{ab} = \mathbf{0}$, i.e. if the Lie derivative with respect to κ^a of the metric vanishes.³⁹ This is equivalent to the requirement that the “flow maps” generated by κ^a are all isometries. (See Wald [1984, p. 441].) For this reason, Killing fields are sometimes called “infinitesimal generators of smooth one-parameter families of isometries” or “infinitesimal symmetries”. The defining condition can also be expressed as⁴⁰

$$(37) \quad \nabla_{(a} \kappa_{b)} = \mathbf{0}.$$

This is “Killing’s equation”.

Given any two smooth vector fields ξ^a and μ^a on M , the *bracket* or *commutator* field $[\xi, \mu]^a$ defined by $[\xi, \mu]^a = \mathcal{L}_\xi \mu^a$ is also smooth. The set of smooth vector fields on M

³⁹Once again, we drop the index on κ here to avoid giving the impression that $\mathcal{L}_\kappa g_{ab}$ is a three index tensor field. A similar remark applies to our bracket notation below.

⁴⁰This follows since $\mathcal{L}_\kappa g_{ab} = \kappa^n \nabla_n g_{ab} + g_{nb} \nabla_a \kappa^n + g_{an} \nabla_b \kappa^n$, and ∇_a is compatible with g_{ab} , i.e. $\nabla_n g_{ab} = \mathbf{0}$.

forms a Lie algebra with respect to this operation, i.e. the bracket operation is linear in each slot; it is anti-symmetric ($[\xi, \mu]^a = -[\mu, \xi]^a$); and it satisfies the Jacobi identity

$$(38) \quad [[\xi, \mu], \nu]^a + [[\nu, \xi], \mu]^a + [[\mu, \nu], \xi]^a = \mathbf{0}$$

for all smooth vector fields ξ^a , μ^a , and ν^a on M . It turns out that the bracket field of two Killing fields is also a Killing field. So it follows, as well, that the set of Killing fields on M has a natural Lie algebra structure.

The discussion of smooth symmetries in spacetime, and their associated conserved quantities, is naturally cast in the language of Killing fields. For example, we can use the latter to capture precisely the following intuitive notions.

- (1) (M, g_{ab}) is *stationary* if it has a Killing field that is everywhere timelike.
- (2) (M, g_{ab}) is *static* if it has a Killing field that is everywhere timelike and locally hypersurface orthogonal.
- (3) (M, g_{ab}) is *homogeneous* if its Killing fields, at every point of M , span the tangent space.
- (4) (M, g_{ab}) is *spatially homogeneous* if it has a smooth unit timelike field ξ^a such that, at every point of M , its Killing fields span the three-dimensional space of vectors orthogonal to ξ^a .
- (5) (M, g_{ab}) is *axially symmetric* if it has a Killing field that (i) is everywhere spacelike, and (ii) has integral curve images that are closed. (The “axis” in this case is the set of points, possibly empty, where the Killing field vanishes.)
- (6) (M, g_{ab}) is *spherically symmetric* if it has three Killing fields $\overset{1}{\sigma}^a, \overset{2}{\sigma}^a, \overset{3}{\sigma}^a$ that (i) are everywhere spacelike, (ii) are linearly dependent at every point, i.e. $\overset{1}{\sigma}^{[a} \overset{2}{\sigma}^b \overset{3}{\sigma}^{c]} = \mathbf{0}$, and (iii) exhibit the same commutation relations as do the generators of the rotation group in three dimensions:

$$(39) \quad [\overset{1}{\sigma}, \overset{2}{\sigma}]^a = \overset{3}{\sigma}^a, \quad [\overset{2}{\sigma}, \overset{3}{\sigma}]^a = \overset{1}{\sigma}^a, \quad [\overset{3}{\sigma}, \overset{1}{\sigma}]^a = \overset{2}{\sigma}^a.$$

The distinction between stationary and static spacetimes should be clear from our discussion in the preceding section. (Recall proposition 9.) Roughly speaking, in a stationary spacetime there is a “timelike flow” that preserves all spacetime distances. But the flow can exhibit rotation. (Think of a perfect whirlpool.) It is the latter possibility that is ruled out when one passes to the definition of a static spacetime.

Now we consider, very briefly, two types of conserved quantity. One is an attribute of massive point particles, the other of extended bodies. Let κ^a be an arbitrary Killing field, and let $\gamma : I \rightarrow M$ be a timelike curve, with unit tangent field ξ^a , whose image is the worldline of a point particle with mass $m > 0$. Consider the quantity $J = (P^a \kappa_a)$, where $P^a = m \xi^a$ is the four-momentum of the particle. It certainly need not be constant on $\gamma[I]$. But it will be if γ is a geodesic. For in that case, $\xi^n \nabla_n \xi^a = \mathbf{0}$ and hence, by (37),

$$(40) \quad \xi^n \nabla_n J = m (\kappa_a \xi^n \nabla_n \xi^a + \xi^n \xi^a \nabla_n \kappa_a) = m \xi^n \xi^a \nabla_{(n} \kappa_{a)} = \mathbf{0}.$$

Thus, the value of J (construed as an attribute of massive point particles) is constant for *free* particles.

We refer to J as the conserved quantity associated with κ^a . If κ^a is timelike, and if the flow maps determined by κ^a have the character of translations⁴¹, then J is called the *energy* of the particle (associated with κ^a).⁴² If it is spacelike, and if the flow maps have the character of translations, then J is called the component of *linear momentum* of the particle (associated with κ^a). Finally, if κ^a is spacelike, and if the flow maps have the character of rotations, then it is called the component of *angular momentum* of the particle (associated with κ^a).

It is useful to keep in mind a certain picture that helps one to “see” why the angular momentum of free particles (to take that example) is conserved. It involves an analogue of angular momentum in Euclidean plane geometry. Figure 4 shows a rotational Killing field κ^a in the Euclidean plane, the image of a geodesic (i.e. a line L), and the tangent field ξ^a to the geodesic. Consider the quantity $J = \xi^a \kappa_a$, i.e. the inner product of ξ^a with κ^a , along L . Exactly the same proof as before (in equation (40)) shows that J is constant along L .⁴³ But here we can better visualize the assertion.

Let us temporarily drop indices and write $\kappa \cdot \xi$ as one would in ordinary Euclidean vector calculus (rather than $\xi^a \kappa_a$). Let p be the point on L that is closest to the center point where κ vanishes. At that point, κ is parallel to ξ . As one moves away from p along L , in either direction, the length $\|\kappa\|$ of κ grows, but the angle $\angle(\kappa, \xi)$ between the vectors increases as well. It is at least plausible from the picture (and easy to check directly with an argument involving similar triangles) that the length of the projection of κ onto the line is constant. Equivalently, the inner product $\kappa \cdot \xi = \cos(\angle(\kappa, \xi)) \|\kappa\| \|\xi\|$ is constant.

That is how to think about the conservation of angular momentum for free particles in relativity theory. It does not matter that in the latter context we are dealing with a Lorentzian metric and allowing for curvature. The claim is still that a certain inner product of vector fields remains constant along a geodesic, and we can still think of that constancy as arising from a compensatory balance of two factors.

Let us now turn to the second type of conserved quantity, the one that is an attribute of extended bodies. Let κ^a be an arbitrary Killing field, and let T_{ab} be the energy-momentum field associated with some matter field. Assume it satisfies the conservation condition. Then $(T^{ab} \kappa_b)$ is divergence free:

$$(41) \quad \nabla_a (T^{ab} \kappa_b) = \kappa_b \nabla_a T^{ab} + T^{ab} \nabla_a \kappa_b = T^{ab} \nabla_{(a} \kappa_{b)} = \mathbf{0}.$$

⁴¹In Minkowski spacetime, one has an unambiguous classification of Killing fields as generators of translations, spatial rotations, boosts (and linear combinations of them). No such classification is available in general. Killing fields are just Killing fields. But sometimes a Killing field in a curved spacetime resembles a Killing field in Minkowski spacetime in certain respects, and then the terminology may carry over naturally. For example, in the case of asymptotically flat spacetimes, one can classify Killing fields by their asymptotic behavior.

⁴²If κ^a is of unit length everywhere, this usage accords well with that in section 2.3. For there ascriptions of energy to point particles were made relative to unit timelike vectors, and the value of the energy at any point was taken to be the inner product of that unit timelike vector with the particle’s four-momentum vector. If κ^a is, at least, of constant length, then one can always rescale it so as to achieve agreement of usage. But, in general, Killing fields, timelike or otherwise, are not of constant length, and so the current usage must be regarded as a generalization of that earlier usage.

⁴³The mass m played no special role.




Figure 4. κ^a is a rotational Killing field. (It is everywhere orthogonal to a circle radius, and proportional to it in length.) ξ^a is a tangent vector field of constant length on the line. The inner-product between them is constant. (Equivalently, the length of the projection of κ^a onto the line is constant.)

(The second equality follows from the conservation condition for T^{ab} (in section 2.4) and the symmetry of T^{ab} ; the third from the fact that κ^a is a Killing field.) It is natural, then, to apply Stokes' theorem to the vector field $(T^{ab}\kappa_b)$.

Consider a bounded system with aggregate energy-momentum field T_{ab} in an otherwise empty universe. Then there exists a (possibly huge) timelike world tube such that T_{ab} vanishes outside the tube (and vanishes on its boundary). Let S_1 and S_2 be (non-intersecting) spacelike hypersurfaces that cut the tube as in figure 5, and let N be the segment of the tube falling between them (with boundaries included). By Stokes' theorem,

$$\begin{aligned} \int_{S_2} (T^{ab}\kappa_b) dS_a - \int_{S_1} (T^{ab}\kappa_b) dS_a &= \int_{S_2 \cap \partial N} (T^{ab}\kappa_b) dS_a - \int_{S_1 \cap \partial N} (T^{ab}\kappa_b) dS_a \\ &= \int_{\partial N} (T^{ab}\kappa_b) dS_a = \int_N \nabla_a (T^{ab}\kappa_b) dV = 0. \end{aligned}$$

Thus, the integral $\int_S (T^{ab}\kappa_b) dS_a$ is independent of the choice of spacelike hypersurface S intersecting the world tube, and is, in this sense, a conserved quantity (construed as an attribute of the system confined to the tube). An "early" intersection yields the same value as a "late" one. Again, the character of the background Killing field κ^a determines our description of the conserved quantity in question. If κ^a is timelike, we take $\int_S (T^{ab}\kappa_b) dS_a$ to be the aggregate energy of the system (associated with κ^a). And so forth.

For further discussion of symmetry and conservation principles in general relativity, see Brading and Castellani (this volume, chapter 13, sections 6 and 7).




Figure 5. The integrated energy (relative to a background timelike Killing field) over the intersection of the world tube with a spacelike hypersurface is independent of the choice of hypersurface.

3 SPECIAL TOPICS

3.1 Relative Simultaneity in Minkowski Spacetime

We noted in section 2.3, when discussing the decomposition of vectors at a point into their “temporal” and “spatial” components relative to a four-velocity vector there, that we were taking for granted the standard identification of relative simultaneity with orthogonality. Here we return to consider the justification of that identification.

Rather than continue to cast the discussion as one concerning the decomposition of the tangent space at a particular point, it is convenient to construe it instead as one about the structure of Minkowski spacetime, the regime of so-called “special relativity”. Doing so will bring it closer to the framework in which traditional discussions of the status of the relative simultaneity relation have been conducted.

Minkowski spacetime is a relativistic spacetime (M, g_{ab}) characterized by three conditions: (i) M is the manifold \mathbb{R}^4 ; (ii) (M, g_{ab}) is flat, i.e. g_{ab} has vanishing Riemann curvature everywhere; and (iii) (M, g_{ab}) is geodesically complete, i.e. every geodesic (with respect to g_{ab}) can be extended to arbitrarily large parameter values in both directions.

By virtue of these conditions, Minkowski spacetime can be identified with its tangent space at any point, and so it inherits the structure of a “metric affine space” in the following sense. Pick any point o in M , and let V be the tangent space M_o at o . Then there is a map $(p, q) \mapsto \vec{pq}$ from $M \times M$ to V with the following two properties.

- (1) For all p, q and r in M , $\vec{pq} + \vec{qr} = \vec{pr}$.
- (2) For all p in M , the induced map $q \mapsto \vec{pq}$ from M to V is a bijection.⁴⁴

⁴⁴If exp is the exponential map from M_o to M , we can take \vec{pq} to be the vector

$$(exp^{-1}(q) - exp^{-1}(p))$$

The triple consisting of the point set M , the vector space V , and the map $(p, q) \mapsto \vec{pq}$ forms an *affine space*. If we add to this triple the inner product on V defined by g_{ab} it becomes a (Lorentzian) *metric affine space*. (For convenience we will temporarily drop the index notation and write $\langle v, w \rangle$ instead of $g_{ab}v^a w^b$ for v and w in V .) We take all this structure for granted in what follows, i.e. we work with Minkowski spacetime and construe it as a metric affine space in the sense described. This will simplify the presentation considerably.

We also use an obvious notation for orthogonality. Given four points p, q, r, s in M , we write $\vec{pq} \perp \vec{rs}$ if $\langle \vec{pq}, \vec{rs} \rangle = 0$. And given a line⁴⁵ L in M , we write $\vec{pq} \perp L$ if $\vec{pq} \perp \vec{rs}$ for all points r and s on L .

Now consider a timelike line L in M . What pairs of points (p, q) in M should qualify as being “simultaneous relative to L ”? That is the question we are considering. The standard answer is that they should do so precisely if $\vec{pq} \perp L$.

In traditional discussions of relative simultaneity, the standard answer is often cast in terms of “epsilon” values. The connection is easy to see. Let p be any point that is not on our timelike line L . Then there exist unique points r and s on L (distinct from one another) such that \vec{rp} and \vec{ps} are future-directed null vectors. (See figure 6.) Now let q




Figure 6. The $\epsilon = \frac{1}{2}$ characterization of relative simultaneity: p and q are simultaneous relative to L iff q is midway between r and s .

be any point on L . (We think of it as a candidate for being judged simultaneous with p

in M_o . All other standard properties of affine spaces follow from these two. E.g. it follows that $\vec{pq} = \mathbf{0} \iff p = q$, for all p and q in M . (Here $\mathbf{0}$ is the zero vector in V .)

⁴⁵In the present context we can characterize a *line* in more than one way. We can take it to be the image of a maximally extended geodesic that is non-trivial, i.e. not a point. Equivalently, we can take it to be a set of points of the form $\{r : \vec{pr} = \epsilon \vec{pq} \text{ for some } \epsilon \in \mathbb{R}\}$ where p and q are any two (distinct) points in M .

relative to L .) Then $\vec{r}\vec{q} = \epsilon \vec{r}\vec{s}$ for some $\epsilon \in \mathbb{R}$. A simple computation⁴⁶ shows that

$$(42) \quad \epsilon = \frac{1}{2} \iff \vec{p}\vec{q} \perp \vec{r}\vec{s}.$$

So the standard (orthogonality) relation of relative simultaneity in special relativity may equally well be described as the “ $\epsilon = \frac{1}{2}$ ” relation of relative simultaneity.

Yet another equivalent formulation involves the “one-way speed of light”. Suppose a light ray travels from r to p with speed c_+ relative to L , and from p to s with speed c_- relative to L . We saw in section 2.3 that if one adopts the standard criterion of relative simultaneity, then it follows that $c_+ = c_-$. (Indeed, in that case, both c_+ and c_- turn out to be 1.) The converse is true as well. For if $c_+ = c_-$, then, as determined relative to L , it should take as much time for light to travel from r to p as from p to s . And in that case, a point q on L should be judged simultaneous with p relative to L precisely if it is midway between r and s . So we are led, once again, to the “ $\epsilon = \frac{1}{2}$ ” relation of relative simultaneity.

Now is adoption of the standard relation a matter of convention, or is it in some significant sense forced on us?

There is, of course, a large literature devoted to this question.⁴⁷ It is not my purpose to review it here, but I do want to draw attention to certain remarks of Howard Stein [1991, pp. 153-4] that seem to me particularly insightful. He makes the point that determinations of conventionality require a context.

There are really two distinct aspects to the issue of the “conventionality” of Einstein’s concept of relative simultaneity. One may assume the position of Einstein himself at the outset of his investigation — that is, of one confronted by a problem, trying to find a theory that will deal with it satisfactorily; or one may assume the position of (for instance) Minkowski — that is, of one confronted with a theory already developed, trying to find its most adequate and instructive formulation.

The problem Einstein confronted was (in part) that of trying to account for our apparent inability to detect any motion of the earth with respect to the “aether”. A crucial element of his solution was the proposal that we think about simultaneity a certain way (i.e. in terms of the “ $\epsilon = \frac{1}{2}$ criterion”), and resolutely follow through on the consequences of doing so. Stein emphasizes just how different that proposal looks when we consider it,

⁴⁶First note that, since $\vec{p}\vec{s}$ and $\vec{p}\vec{r}$ are null,

$$0 = \langle \vec{p}\vec{s}, \vec{p}\vec{s} \rangle = \langle \vec{p}\vec{r} + \vec{r}\vec{s}, \vec{p}\vec{r} + \vec{r}\vec{s} \rangle = 2\langle \vec{p}\vec{r}, \vec{r}\vec{s} \rangle + \langle \vec{r}\vec{s}, \vec{r}\vec{s} \rangle.$$

It follows that

$$\langle \vec{p}\vec{q}, \vec{r}\vec{s} \rangle = \langle \vec{p}\vec{r} + \vec{r}\vec{q}, \vec{r}\vec{s} \rangle = \langle \vec{p}\vec{r} + \epsilon \vec{r}\vec{s}, \vec{r}\vec{s} \rangle = \langle \vec{p}\vec{r}, \vec{r}\vec{s} \rangle + \epsilon \langle \vec{r}\vec{s}, \vec{r}\vec{s} \rangle = (\epsilon - \frac{1}{2}) \langle \vec{r}\vec{s}, \vec{r}\vec{s} \rangle,$$

which implies (42).

⁴⁷Classic statements of the conventionalist position can be found in Reichenbach [1958] and Grünbaum [1973]. Grünbaum has recently responded to criticism of his views in [forthcoming]. An overview of the debate with many references can be found in Janis [2002].

not from Einstein's initial position, but rather from the vantage point of the finished theory, i.e. relativity theory conceived as an account of invariant spacetime structure.

[For] Einstein, the question (much discussed since Reichenbach) whether the evidence really shows that the speed of light *must* be regarded as the same in all directions and for all observers is not altogether appropriate. A person devising a theory does not have the responsibility, at the outset, of showing that the theory being developed is the only possible one given the evidence. [But] once Einstein's theory had been developed, and had proved successful in dealing with all relevant phenomena, the case was quite transformed; for we know that *within* this theory, there is only one "reasonable" concept of simultaneity (and in terms of that concept, the velocity of light is indeed as Einstein supposed); therefore an alternative will only present itself if someone succeeds in constructing, not simply a different empirical criterion of simultaneity, but an essentially different (and yet viable) theory of electrodynamics of systems in motion. No serious alternative theory is in fact known. (emphasis in original)

My goal in the remainder of this section is to formulate three elementary uniqueness results, closely related to one another, that capture the sense in which "there is only one 'reasonable' concept of (relative) simultaneity" within the framework of Minkowski spacetime.

It will help to first consider an analogy. In some formulations of Euclidean plane geometry, the relation of congruence between angles is taken as primitive along with that of congruence between line segments (and other relations suitable for formulating axioms about affine structure). But suppose we have a formulation in which it is not, and we undertake to *define* a notion of angle-congruence in terms of the other primitives. The standard angle-congruence relation can certainly be defined this way, and there is a clear sense in which it is the only reasonable candidate. Consider any two angles in the Euclidean plane. (Let's agree that an "angle" consists of two rays, i.e. half-lines, with a common initial point.) Whatever else is the case, presumably, it is only reasonable to count them as congruent, i.e. equal in "size", if there is an isometry of the Euclidean plane that maps one angle onto the other.⁴⁸ So though we have here a notion of angle-congruence that is introduced "by definition", there is no interesting sense in which it is conventional in character.

A situation very much like this arises if we think about "one-way light speeds" in terms of Minkowskian spacetime geometry. Indeed, the claim that the speed of light *in vacuo* is the same in all directions and for all inertial observers is naturally represented as a claim about angle congruence (for a special type of angle) in Minkowski spacetime.

Let us take a "light-speed angle" to be a triple of the form (p, T, N) , where p is a point in M , T is a future-pointing timelike ray with initial point p , and N is a future-pointing null ray with initial point p . (See figure 7.)

⁴⁸In this context, a one-to-one map of the Euclidean plane onto itself is an "isometry" if it preserves the relation of congruence between line segments.




Figure 7. Congruent “light speed angles” in Minkowski spacetime.

Then we can represent systematic attributions of one-way light speed as maps of the form: $(p, T, N) \mapsto v(p, T, N)$. (We understand $v(p, T, N)$ to be the speed that an observer with (half) worldline T at p assigns to the light signal with (half) worldline N .) So, for example, the principle that the speed of light is the same in all directions and for all inertial observers comes out as the condition that $v(p, T, N) = v(p', T', N')$ for all light-speed angles (p, T, N) and (p', T', N') .

Now it is natural to regard $v(p, T, N)$ as a measure of the “size” of the angle (p, T, N) . If we do so, then, just as in the Euclidean case, we can look to the background metric to decide when two angles have the same size. That is, we can take them to be congruent iff there is an isometry of Minkowski spacetime that maps one to the other. But on this criterion, all light-speed angles are congruent (proposition 10). So we are led back to the principle that the (one-way) speed of light is the same in all directions and for all inertial observers and, hence, back to the standard relative simultaneity relation.

PROPOSITION 10. *Let (p, T, N) and (p', T', N') be any two light speed angles in Minkowski spacetime. Then there is an isometry ϕ of Minkowski spacetime such that $\phi(p) = p'$, $\phi[T] = T'$, and $\phi[N] = N'$.⁴⁹*

Once again, let L be a timelike line in M , and let Sim_L be the standard relation of simultaneity relative to L . (So $(p, q) \in Sim_L$ iff $\vec{pq} \perp L$, for all p and q in M .) Further, let S be an arbitrary two-place relation on M that we regard as a candidate for the relation of “simultaneity relative to L ”. Our second uniqueness result asserts that if S satisfies three conditions, including an invariance condition, then $S = Sim_L$.⁵⁰

The first two conditions are straightforward.

- (S1) S is an equivalence relation (i.e. S is reflexive, symmetric, and transitive).
- (S2) For all points $p \in M$, there is a unique point $q \in L$ such that $(p, q) \in S$.

If S satisfies (S1), it has an associated family of equivalence classes. We can think of them as “simultaneity slices” (as determined relative to L). Then (S2) asserts that every

⁴⁹The required isometry can be realized in the form $\phi = \varphi_3 \circ \varphi_2 \circ \varphi_1$ where (i) φ_1 is a translation that takes p to p' , (ii) φ_2 is a boost (based at p') that maps $\varphi_1[T]$ to T' , and (iii) φ_3 is a rotation about T' that maps $(\varphi_2 \circ \varphi_1)[N]$ to N' .

⁵⁰Many propositions of this form can be found in the literature. (See Budden [1998] for a review.) Ours is intended only as an example. There are lots of possibilities here depending on exactly how one formulates the conditions that S must satisfy. The proofs are all very much the same.

simultaneity slice intersects L in exactly one point. Note that if $S = Sim_L$, then (S1) and (S2) are satisfied. For in this case, the equivalence classes associated with S are hyperplanes orthogonal to L , and these clearly intersect L in exactly one point.

The third, invariance condition is intended to capture the requirement that S is determined by, or definable in terms of, the background geometric structure of Minkowski spacetime and by L itself. The one subtle point here is whether temporal orientation is taken to count as part of that background geometric structure or not. Let's assume for the moment that it does not.

Let $\varphi: M \rightarrow M$ be an isometry of (M, g_{ab}) . We say it is an L -isometry if, in addition, it preserves L , i.e. if, for all points p in M , $p \in L \iff \varphi(p) \in L$. The set of L -isometries is generated by maps of the following three types: (a) translations ("up" and "down") in the direction of L , (b) spatial rotations that leave fixed every point in L , and (c) temporal reflections with respect to spacelike hyperplanes orthogonal to L . We will say that our two-place relation S is L -invariant if it preserved under all L -isometries, i.e. if for all L -isometries $\varphi: M \rightarrow M$, and all points $p, q \in M$,

$$(43) \quad (p, q) \in S \iff (\varphi(p), \varphi(q)) \in S.$$

We can now formulate the second uniqueness result.⁵¹

PROPOSITION 11. *Let L be a timelike line, and let S be a two-place relation on M that satisfies conditions (S1) and (S2), and is L -invariant. Then $S = Sim_L$.*

As it turns out, the full strength of L -invariance is not needed here. It suffices to require that S is preserved under all L -isometries of type (c).⁵²

Suppose now that we *do* want to consider temporal orientation as part of the background structure that may play a role in the determination of S . Then we need to recast the invariance condition. Let us say that an L -isometry $\varphi: M \rightarrow M$ is an (L, \uparrow) -isometry if it (also) preserves temporal orientation, i.e. if for all timelike vectors \vec{p}, \vec{q} , $\overrightarrow{\varphi(p)}\overrightarrow{\varphi(q)}$ is co-oriented with $\vec{p}\vec{q}$. And let us say that S is (L, \uparrow) -invariant if it is preserved under all (L, \uparrow) -isometries. (So, to be (L, \uparrow) -invariant, S must be preserved under all L -isometries of type (a) and (b), but need not be preserved under those of type (c).)

(L, \uparrow) -invariance is a weaker condition than L -invariance and, in fact, is too weak to deliver the uniqueness result we want. It turns out there are many two-place relations S on M other than Sim_L that satisfy (S1), (S2), and are (L, \uparrow) -invariant. Their associated "simultaneity slices" are "flat cones" that are preserved under L -isometries of type (a) and (b), but not (c). (See figure 8.)

But we can still get a uniqueness result if we change the set-up slightly, and think of simultaneity as determined, not relative to individual timelike lines, but, rather, relative to families of parallel timelike lines. Let us officially take a *frame* in M to be a set of

⁵¹It is a close variant of one presented in Hogarth [1993].

⁵²The key step in the proof is the following. Let p be a point in M . By (S2), there is a unique point q on L such that $(p, q) \in S$. Let $\varphi: M \rightarrow M$ be a reflection with respect to the hyperplane orthogonal to L that passes through p . Then $\varphi(p) = p$, $\varphi(q) \in L$, and S is invariant under φ . Hence $(p, \varphi(q)) = (\varphi(p), \varphi(q)) \in S$. Since $\varphi(q) \in L$, it follows by the uniqueness condition in (S2) that $\varphi(q) = q$. But the only points left fixed by φ are those on the hyperplane orthogonal to L that passes through p . So p and q are both on that hyperplane, and $\vec{p}\vec{q}$ is orthogonal to L , i.e. $(p, q) \in Sim_L$.




Figure 8. The “flat cones” displayed are the simultaneity slices associated with a two-place relation S that satisfies conditions (S1) and (S2), and is (L, \uparrow) -invariant (but not L -invariant).

parallel timelike lines \mathcal{L} that is maximal in the sense that every point in M falls on one (and only one) of them. With hardly any work, we can recast our previous notions in terms of frames rather than lines.

In what follows, let \mathcal{L} be some fixed frame. Given any two lines L and L' in \mathcal{L} , $Sim_L = Sim_{L'}$. (Since L and L' are parallel, any vector orthogonal to one must be orthogonal to the other.) So we can, without ambiguity, make reference to $Sim_{\mathcal{L}}$ (the standard relation of simultaneity relative to \mathcal{L}). Let $\varphi : M \rightarrow M$ be an isometry of (M, g_{ab}) . We say it is an \mathcal{L} -isometry if, for all L in \mathcal{L} , the line $\varphi[L]$ is also in \mathcal{L} . And we say that it is an (\mathcal{L}, \uparrow) -isometry if, in addition, it preserves temporal orientation.

If L is a line in \mathcal{L} , then the set of \mathcal{L} -isometries certainly includes all L -isometries of types (a), (b), and (c) above. But it includes, in addition, (d) translations taking L to some other line in \mathcal{L} , and (e) isometries that leave fixed the points on some line in \mathcal{L} other than L . If we restrict attention to (\mathcal{L}, \uparrow) -isometries, we lose maps of type (c), but we retain those of types (a), (b), (d), and (e). Invariance under this larger class is sufficient to drive a uniqueness result.

We say (of course) that S is \mathcal{L} -invariant if it is preserved under all \mathcal{L} -isometries, and (\mathcal{L}, \uparrow) -invariant if it is preserved under all (\mathcal{L}, \uparrow) -isometries. Our third uniqueness result comes out as follows.⁵³

PROPOSITION 12. *Let \mathcal{L} be a frame, and let S be a two-place on M . Suppose S satisfies (S1) and, for some L in \mathcal{L} , satisfies (S2). Further, suppose S is (\mathcal{L}, \uparrow) -invariant. Then $S = Sim_{\mathcal{L}}$.*

The move from proposition 11 to proposition 12 involves a trade-off. We drop the requirement that S be invariant under maps of type (c), but add the requirement that it be invariant under those of type (d) and (e). (It is a good exercise to check that one does not need the full strength of (\mathcal{L}, \uparrow) -invariance in proposition 12. It suffices to require that S

⁵³It is closely related to propositions in Spirtes [1981], Stein [1991], and Budden [1998].

be invariant under maps of type (a) and (e), or, alternatively, invariant under maps of type (b) and (d).)

Once again, many variations of these results can be found in the literature. For example, if one subscribes to a “causal theory of time (or spacetime)”, one will want to consider what candidate simultaneity relations are determined by the causal structure of Minkowski spacetime (in addition to the line L). Let $\varphi: M \rightarrow M$ be a bijection. We say it is a *causal isomorphism* if it preserves the symmetric relation of causal connectivity, i.e. if, for all points p and q in M , \overrightarrow{pq} is a causal vector iff $\overrightarrow{\varphi(p)\varphi(q)}$ is one. We say it is *L -causal isomorphism*, an *(L, \uparrow) -causal isomorphism*, and *\mathcal{L} -causal isomorphism*, or an *(\mathcal{L}, \uparrow) -causal isomorphism* if, in each case, it preserves the specified additional structure. Since isometries of Minkowski spacetime automatically qualify as causal isomorphisms, we can replace the invariance condition in proposition 11 with the requirement that S be invariant under all L -causal isomorphisms; and we can replace the one in proposition 12 with the requirement that S be invariant under all (\mathcal{L}, \uparrow) -causal isomorphisms.

3.2 Geometrized Newtonian Gravitation Theory

The “geometrized” formulation of Newtonian gravitation theory was first introduced by Cartan [1923; 1924], and Friedrichs [1927], and later developed by Trautman [1965], Künzle [1972; 1976], Ehlers [1981], and others.

It is significant for several reasons. (1) It shows that several features of relativity theory once thought to be uniquely characteristic of it do not distinguish it from (a suitably reformulated version of) Newtonian gravitation theory. The latter too can be cast as a “generally covariant” theory in which (a) gravity emerges as a manifestation of spacetime curvature, and (b) spacetime structure is “dynamical”, i.e. participates in the unfolding of physics rather than being a fixed backdrop against which it unfolds.

(2) It helps one to see where Einstein’s equation “comes from”, at least in the empty-space case. (Recall the discussion in section 2.5.) It also allows one to make precise, in coordinate-free, geometric language, the standard claim that Newtonian gravitation theory (or, at least, a certain generalized version of it) is the “classical limit” of general relativity. (See Künzle [1976] and Ehlers [1981].)

(3) It clarifies the gauge status of the Newtonian gravitational potential. In the geometrized formulation of Newtonian theory, one works with a single curved derivative operator $\overset{g}{\nabla}_a$. It can be decomposed (in a sense) into two pieces — a flat derivative operator ∇_a and a gravitational potential ϕ — to recover the standard formulation of the theory.⁵⁴ But in the absence of special boundary conditions, the decomposition will not be unique. Physically, there is no unique way to divide into “inertial” and “gravitational” components the forces experienced by particles. Neither has any direct physical significance. Only their “sum” does. It is an attractive feature of the geometrized formulation that it trades in two gauge quantities for this sum.

(4) The clarification described in (3) also leads to a solution, or dissolution, of an old

⁵⁴As understood here, the “standard” formulation is not that found in textbooks, but rather a “generally covariant” theory of four-dimensional spacetime structure in which gravity is not geometrized.

conceptual problem about Newtonian gravitation theory, namely the apparent breakdown of the theory when applied (in cosmology) to a hypothetically infinite, homogeneous mass distribution. (See Malament [1995] and Norton [1995; 1999].)

In what follows, we give a brief overview of the geometrized formulation of Newtonian gravitation theory, and say a bit more about points (1) and (3). We start by characterizing a new class of geometrical models for the spacetime structure of our universe (or subregions thereof) that is broad enough to include the models considered in both the standard and geometrized versions of Newtonian theory. We take a *classical spacetime* to be a structure $(M, t_{ab}, h^{ab}, \nabla_a)$ where (i) M is a smooth, connected, four-dimensional differentiable manifold; (ii) t_{ab} is a smooth, symmetric, covariant tensor field on M of signature $(1, 0, 0, 0)$;⁵⁵ (iii) h^{ab} is a smooth, symmetric, contravariant tensor field on M of signature $(0, 1, 1, 1)$; (iv) ∇_a is a smooth derivative operator on M ; and (v) the following two conditions are met:

$$(44) \quad h^{ab} t_{bc} = \mathbf{0}$$

$$(45) \quad \nabla_a t_{bc} = \mathbf{0} = \nabla_a h^{bc}.$$

We refer to them, respectively, as the “orthogonality” and “compatibility” conditions.

M is interpreted as the manifold of point events (as before); t_{ab} and h^{ab} are understood to be temporal and spatial metrics on M , respectively. Collectively, the objects t_{ab} , h^{ab} , and ∇_a represent the spacetime structure presupposed by classical, Galilean relativistic dynamics. We review, briefly, how they do so.

In what follows, let $(M, t_{ab}, h^{ab}, \nabla_a)$ be a fixed classical spacetime.

Consider, first, t_{ab} . Given any vector ξ^a at a point, it assigns a “temporal length” $(t_{ab} \xi^a \xi^b)^{\frac{1}{2}} \geq 0$. The vector ξ^a is classified as *timelike* or *spacelike* depending on whether its temporal length is positive or zero. It follows from the signature of t_{ab} that the subspace of spacelike vectors at any point is three-dimensional. It also follows from the signature that at every point there exists a covariant vector t_a , unique up to sign, such that $t_{ab} = t_a t_b$. We say that the structure $(M, t_{ab}, h^{ab}, \nabla_a)$ is *temporally orientable* if there is a continuous (globally defined) vector field t_a such that this decomposition holds at every point. Each such field t_a (which, in fact, must be smooth because t_{ab} is) determines a *temporal orientation*. A timelike vector ξ^a qualifies as *future-directed* relative to t_a if $t_a \xi^a > 0$; otherwise it is *past-directed*. Let us assume in what follows that $(M, t_{ab}, h^{ab}, \nabla_a)$ is temporally orientable and that a temporal orientation t_a has been selected.

⁵⁵The signature condition for t_{ab} is equivalent to the requirement that, at every point in M , the tangent space there have a basis $\overset{1}{\xi}^a, \dots, \overset{4}{\xi}^a$ such that, for all i and j in $\{1, 2, 3, 4\}$, $t_{ab} \overset{i}{\xi}^a \overset{j}{\xi}^b = 0$ if $i \neq j$, and

$$t_{ab} \overset{i}{\xi}^a \overset{i}{\xi}^b = \begin{cases} 1 & \text{if } i = 1 \\ 0 & \text{if } i = 2, 3, 4. \end{cases}$$

Similarly, the signature condition for h^{ab} in (iii) requires that, at every point, the cotangent space there have a basis $\overset{1}{\sigma}_a, \dots, \overset{4}{\sigma}_a$ such that, for all i and j in $\{1, 2, 3, 4\}$, $h^{ab} \overset{i}{\sigma}_a \overset{j}{\sigma}_b = 0$ if $i \neq j$, and

$$h^{ab} \overset{i}{\sigma}_a \overset{i}{\sigma}_b = \begin{cases} 0 & \text{if } i = 1 \\ 1 & \text{if } i = 2, 3, 4. \end{cases}$$

From the compatibility condition, it follows that t_a is closed, i.e. $\nabla_{[a} t_{b]} = 0$. So, at least locally, it must be exact, i.e. of form $t_a = \nabla_a t$ for some smooth function t . We call any such function a *time function*. If M has a suitable global structure, e.g. if it is simply connected, then a globally defined time function $t : M \rightarrow \mathbb{R}$ must exist. In this case, spacetime can be decomposed into a one-parameter family of global ($t = \text{constant}$) “time slices”. One can speak of “space” at a given “time”. A different choice of time function would result in a different zero-point for the time scale, but would induce the same time slices and the same elapsed intervals between them.

We say that a smooth curve is *timelike* (respectively *spacelike*) if its tangent field is timelike (respectively spacelike) at every point. In what follows, unless indication is given to the contrary, it should further be understood that a “timelike curve” is future-directed and parametrized by its t_{ab} -length. In this case, its tangent field ξ^a satisfies the normalization condition $t_a \xi^a = 1$. Also, in this case, if a particle happens to have the image of the curve as its worldline, then, at any point, ξ^a is called the particle’s *four-velocity*, and $\xi^n \nabla_n \xi^a$ its *four-acceleration*, there.⁵⁶ If the particle has mass m , then its four-acceleration field satisfies the equation of motion

$$(46) \quad F^a = m \xi^n \nabla_n \xi^a,$$

where F^a is a spacelike vector field (on the image of its worldline) that represents the net force acting on the particle. This is, once again, our version of Newton’s second law of motion. Recall (15). Note that the equation makes geometric sense because four-acceleration vectors are necessarily spacelike.⁵⁷

Now consider h^{ab} . It serves as a spatial metric, but just how it does so is a bit tricky. In Galilean relativistic mechanics, we have no notion of spatial length for timelike vectors, e.g. four-velocity vectors, since having one is tantamount to a notion of absolute rest. (We can take a particle to be at rest if its four-velocity has spatial length 0 everywhere.) But we *do* have a notion of spatial length for spacelike vectors, e.g. four-acceleration vectors. (We can, for example, use measuring rods to determine distances between simultaneous events.) h^{ab} serves to give us one without the other.

We cannot take the spatial length of a vector σ^a to be $(h_{ab} \sigma^a \sigma^b)^{\frac{1}{2}}$ because the latter is not well-defined. (Since h^{ab} has degenerate signature, it is not invertible, i.e. there does not *exist* a field h_{ab} satisfying $h^{ab} h_{bc} = \delta^a_c$.) But if σ^a is spacelike, we can use h^{ab} to assign a spatial length to it indirectly. It turns out that: (i) a vector σ^a is spacelike iff it can be expressed in the form $\sigma^a = h^{ab} \lambda_b$, and (ii) if it can be so expressed, the quantity $(h^{ab} \lambda_a \lambda_b)$ is independent of the choice of λ_a . Furthermore, the signature of h^{ab} guarantees that $(h^{ab} \lambda_a \lambda_b) \geq 0$. So if σ^a is spacelike, we can take its spatial length to be $(h^{ab} \lambda_a \lambda_b)^{\frac{1}{2}}$, for any choice of corresponding λ_a .

⁵⁶Here we take for granted an interpretive principle that corresponds to $C1$: (i) a curve is timelike iff its image could be the worldline of a point particle. Other principles we can formulate at this stage correspond to $P1$ and $P2$: (ii) a timelike curve can be reparametrized so as to be a geodesic (with respect to ∇_a) iff its image could be the worldline of a free particle; (iii) clocks record the passage of elapsed t_{ab} -length along their worldlines. (Here, in contrast to the relativistic setting, we have only massive particles to consider; and until we geometrize Newtonian gravity, we do not count a particle as “free” if it is subject to “gravitational force”.)

⁵⁷By the compatibility condition, $t_a \xi^n \nabla_n \xi^a = \xi^n \nabla_n (t_a \xi^a) = \xi^n \nabla_n (1) = 0$.

One final preliminary remark about classical spacetimes is needed. It is crucial for our purposes, as will be clear, that the compatibility condition (45) does not determine a unique derivative operator. (It is a fundamental result that the compatibility condition $\nabla_a g_{bc} = \mathbf{0}$ determines a unique derivative operator if g_{ab} is a semi-Riemannian metric, i.e. a smooth, symmetric field that is invertible (i.e. non-degenerate). But neither t_{ab} nor h^{ab} is invertible.)

Because h^{ab} is not invertible, we cannot raise *and* lower indices with it. But we can, at least, raise indices with it, and it is sometimes convenient to do so. So, for example, if R^a_{bcd} is the Riemann curvature tensor field associated with ∇_a , we can understand R^{ab}_{cd} to be an abbreviation for $h^{bn} R^a_{ncd}$.

Let us now, finally, consider Newtonian gravitation theory. In the standard (non-geometrized) version, one works with a flat derivative operator ∇_a and a gravitational potential ϕ , the latter understood to be a smooth, real-valued function on M . The gravitational force on a point particle with mass m is given by $-m h^{ab} \nabla_b \phi$. (Notice that this is a spacelike vector by the orthogonality condition.) Using our convention for raising indices, we can also express the vector as: $-m \nabla^a \phi$. It follows that if the particle is subject to no forces except gravity, and if it has four-velocity ξ^a , it satisfies the equation of motion

$$(47) \quad -\nabla^a \phi = \xi^n \nabla_n \xi^a.$$

(Here we have just used $-m \nabla^a \phi$ for the left side of (46).) It is also assumed that ϕ satisfies Poisson's equation:

$$(48) \quad \nabla^a \nabla_a \phi = 4 \pi \rho,$$

where ρ is the Newtonian mass-density function (another smooth real-valued function on M). (The expression on the left side is an abbreviation for: $h^{ab} \nabla_b \nabla_a \phi$.)

In the geometrized formulation of the theory, gravitation is no longer conceived as a fundamental "force" in the world, but rather as a manifestation of spacetime curvature (just as in relativity theory). Rather than thinking of point particles as being deflected from their natural straight (i.e. geodesic) trajectories, one thinks of them as traversing geodesics in curved spacetime. So we have a geometry problem. Starting with the structure $(M, t_{ab}, h^{ab}, \nabla_a)$, can we find a new derivative operator $\overset{g}{\nabla}_a$, also compatible with the metrics t_{ab} and h^{ab} , such that a timelike curve satisfies the equation of motion (47) with respect to the original derivative operator ∇_a iff it is a geodesic with respect to $\overset{g}{\nabla}_a$? The following proposition (essentially due to Trautman [1965]) asserts that there is exactly one such $\overset{g}{\nabla}_a$. It also records several facts about the Riemann curvature tensor field $\overset{g}{R}^a_{bcd}$ associated with $\overset{g}{\nabla}_a$.

In formulating the proposition, we make use of the following basic fact about derivative operators. Given any two such operators $\overset{1}{\nabla}_a$ and $\overset{2}{\nabla}_a$ on M , there is a unique smooth tensor field C^a_{bc} , symmetric in its covariant indices, such that, for all smooth fields $\alpha^{a\dots b}_{c\dots d}$

on M ,

$$(49) \quad \overset{2}{\nabla}_n \alpha^{a\dots b}{}_{c\dots d} = \overset{1}{\nabla}_n \alpha^{a\dots b}{}_{c\dots d} + C^r{}_{nc} \alpha^{a\dots b}{}_{r\dots d} + \dots + C^r{}_{nd} \alpha^{a\dots b}{}_{c\dots r} - C^a{}_{nr} \alpha^{r\dots b}{}_{c\dots d} - \dots - C^b{}_{nr} \alpha^{a\dots r}{}_{c\dots d}.$$

In this case, we say that “the action of $\overset{2}{\nabla}_a$ relative to that of $\overset{1}{\nabla}_a$ is given by $C^a{}_{bc}$ ”.⁵⁸ Conversely, given any one derivative operator $\overset{1}{\nabla}_a$ on M , and any smooth, symmetric field $C^a{}_{bc}$ on M , (49) defines a new derivative operator $\overset{2}{\nabla}_a$ on M . (See Wald [1984, p. 33].)

PROPOSITION 13 (Geometrization Theorem). *Let $(M, t_{ab}, h^{ab}, \nabla_a)$ be a classical spacetime with ∇_a flat ($R^a{}_{bcd} = \mathbf{0}$). Further, let ϕ and ρ be smooth real valued functions on M satisfying Poisson’s equation: $\nabla^a \nabla_a \phi = 4\pi \rho$. Finally, let $\overset{g}{\nabla}_a$ be the derivative operator on M whose action relative to that of ∇_a is given by $C^a{}_{bc} = -t_{bc} \nabla^a \phi$. Then all the following hold.*

(G1) $(M, t_{ab}, h^{ab}, \overset{g}{\nabla}_a)$ is a classical spacetime.

(G2) $\overset{g}{\nabla}_a$ is the unique derivative operator on M such that, for all timelike curves on M with four-velocity fields ξ^a ,

$$(50) \quad \xi^n \overset{g}{\nabla}_n \xi^a = \mathbf{0} \iff -\nabla^a \phi = \xi^n \nabla_n \xi^a.$$

(G3) The curvature field $\overset{g}{R}^a{}_{bcd}$ associated with $\overset{g}{\nabla}_a$ satisfies:

$$\overset{g}{R}_{bc} = 4\pi \rho t_{bc} \tag{51}$$

$$\overset{g}{R}{}^{ab}{}_{cd} = \mathbf{0} \tag{52}$$

$$\overset{g}{R}{}^{[a}{}_{(b}{}^{c]}{}_{d)} = \mathbf{0}. \tag{53}$$

(51) is the geometrized version of Poisson’s equation. The proof proceeds by more-or-less straight forward computation using (49).⁵⁹

⁵⁸Clearly, if the action of $\overset{2}{\nabla}_a$ relative to that of $\overset{1}{\nabla}_a$ is given by $C^a{}_{bc}$, then, conversely, the action of $\overset{1}{\nabla}_a$ relative to that of $\overset{2}{\nabla}_a$ is given by $-C^a{}_{bc}$. In the sum on the right side of (49), there is one term involving $C^a{}_{bc}$ for each index in $\alpha^{a\dots b}{}_{c\dots d}$. In each case, the index in question is contracted with $C^a{}_{bc}$, and the term carries a coefficient of +1 or -1 depending on whether the index in question is in covariant (down) or contravariant (up) position. (The components of $C^a{}_{bc}$ in a particular coordinate system are obtained by subtracting the Christoffel symbols associated with $\overset{1}{\nabla}_a$ (in that coordinate system) from those associated with $\overset{2}{\nabla}_a$.)

⁵⁹Here is a sketch. By (49),

$$\overset{g}{\nabla}_a t_{bc} = \nabla_a t_{bc} + C^r{}_{ab} t_{rc} + C^r{}_{ac} t_{br} = \nabla_a t_{bc} + (-t_{ab} \nabla^r \phi) t_{rc} + (-t_{ac} \nabla^r \phi) t_{br}.$$

The first term in the far right sum vanishes by the compatibility condition (45); the second and third do so by the orthogonality condition (44) since, for example, $(\nabla^r \phi) t_{br} = (h^{rm} t_{br}) \nabla_m \phi$. So $\overset{g}{\nabla}_a$ is compatible with t_{bc} . Much the same argument shows that it is also compatible with h^{ab} . This gives us (G1).

We can also work in the opposite direction. In geometrized Newtonian gravitation theory, one *starts* with a curved derivative operator $\overset{g}{\nabla}_a$ satisfying (51), (52), (53), and with the principle that point particles subject to no forces (except “gravity”) traverse geodesics with respect to $\overset{g}{\nabla}_a$. (52) and (53) function as integrability conditions that ensure the possibility of working backwards to recover the standard formulation in terms of a gravitational potential and flat derivative operator.⁶⁰ We have the following recovery, or degeometrization, theorem (also essentially due to Trautman [1965]).

PROPOSITION 14 (Recovery Theorem). *Let $(M, t_{ab}, h^{ab}, \overset{g}{\nabla}_a)$ be a classical spacetime that, together with a smooth, real-valued function ρ on M , satisfies conditions (51), (52), (53). Then, at least locally (and globally if M is, for example, simply connected), there exists a smooth, real-valued function ϕ on M and a flat derivative operator ∇_a such that all the following hold.*

(R1) $(M, t_{ab}, h^{ab}, \nabla_a)$ is a classical spacetime.

(R2) For all timelike curves on M with four-velocity fields ξ^a , the geometrization condition (50) is satisfied.

For (G2), let $\overset{g}{\nabla}_a$ (temporarily) be an arbitrary derivative operator on M whose action relative to that ∇_a is given by some field C^a_{bc} . Let p be an arbitrary point in M , and let ξ^a be the four-velocity field of an arbitrary timelike curve through p . Then, by (49),

$$\xi^n \overset{g}{\nabla}_n \xi^a = \xi^n \nabla_n \xi^a - C^a_{rn} \xi^r \xi^n.$$

It follows that $\overset{g}{\nabla}_n$ will satisfy (G2) iff $C^a_{rn} \xi^r \xi^n = -\nabla^a \phi$ or, equivalently,

$$[C^a_{rn} + (\nabla^a \phi) t_{rn}] \xi^r \xi^n = \mathbf{0},$$

for all future-directed unit timelike vectors ξ^a at all points p . But the space of future-directed unit timelike vectors at any p spans the tangent space M_p there, and the field in brackets is symmetric in its covariant indices.

So, $\overset{g}{\nabla}_n$ will satisfy (G2) iff $C^a_{rn} = -(\nabla^a \phi) t_{rn}$ everywhere.

Finally, for (G3) we use the fact that $\overset{g}{R}^a_{bcd}$ can be expressed as a sum of terms involving R^a_{bcd} and C^a_{bc} (see Wald [1984, p. 184]), and then substitute for C^a_{bc} :

$$\begin{aligned} \overset{g}{R}^a_{bcd} &= R^a_{bcd} + 2 \nabla_{[c} C^a_{d]b} + 2 C^m_{b[c} C^a_{d]n} \\ &= R^a_{bcd} - 2 t_{b[d} \nabla_{c]} \nabla^a \phi = -2 t_{b[d} \nabla_{c]} \nabla^a \phi. \end{aligned}$$

(Here $C^n_{b[c} C^a_{d]n}$ turns out to be $\mathbf{0}$ by the orthogonality condition, and $\nabla_{[c} C^a_{d]b}$ turns out to be $-t_{b[d} \nabla_{c]} \nabla^a \phi$ by the compatibility condition. For the final equality we use our assumption that $R^a_{bcd} = \mathbf{0}$.) (52) and (53) now follow from the orthogonality condition and (for (53)) from the fact that $\nabla^{[c} \nabla^a] \phi = \mathbf{0}$ for any smooth function ϕ . Contraction on ‘a’ and ‘d’ yields

$$\overset{g}{R}_{bc} = t_{bc} (\nabla_a \nabla^a \phi).$$


So (51) follows from our assumption that $\nabla^a \nabla_a \phi = 4 \pi \rho$ (and the fact that $\nabla_a \nabla^a \phi = \nabla^a \nabla_a \phi$).

⁶⁰I am deliberately passing over some subtleties here. Geometrized Newtonian gravitation theory comes in several variant formulations. (See Bain [2004] for a careful review of the differences.) The one presented here is essentially that of Trautman [1965]. In other weaker formulations (such as that in Künzle [1972]), condition (52) is dropped, and it is *not* possible to fully work back to the standard formulation (in terms of a gravitational potential and flat derivative operator) unless special global conditions on spacetime structure are satisfied.

(R3) ∇_a satisfies Poisson's equation: $\nabla^a \nabla_a \phi = 4 \pi \rho$.

The theorem is an existential assertion of this form: given $\overset{g}{\nabla}_a$ satisfying certain conditions, there exists (at least locally) a smooth function ϕ on M and a flat derivative operator ∇_a such that $\overset{g}{\nabla}_a$ arises as the "geometrization" of the pair (∇_a, ϕ) . But, as claimed at the beginning of this section, we do *not* have uniqueness unless special boundary conditions are imposed on ϕ .

For suppose ∇_a is flat, and the pair (∇_a, ϕ) satisfies (R1), (R2), (R3). Let ψ be any smooth function (with the same domain as ϕ) such that $\nabla^a \nabla^b \psi$ vanishes everywhere, but $\nabla^b \psi$ does not.⁶¹ If we set $\bar{\phi} = \phi + \psi$, and take $\bar{\nabla}_a$ to be the derivative operator relative to which the action of $\overset{g}{\nabla}_a$ is given by $\bar{C}^a_{bc} = -t_{bc} \nabla^a \bar{\phi}$, then $\bar{\nabla}_a$ is flat and the pair $(\bar{\nabla}_a, \bar{\phi})$ satisfies conditions (R1), (R2), (R3) as well.⁶²



But, because $\nabla^b \psi$ is non-vanishing (somewhere or other), the pairs (∇_a, ϕ) and $(\bar{\nabla}_a, \bar{\phi})$ are distinct decompositions of $\overset{g}{\nabla}_a$. Relative to the first, a point particle (with mass m and four-velocity ξ^a) has acceleration $\xi^n \nabla_n \xi^a$ and is subject to a gravitational force

⁶¹We can think of $\nabla^b \psi$ as the "spatial gradient" of ψ . The stated conditions impose the requirement that $\nabla^b \psi$ be constant on all spacelike submanifolds ("time slices"), but not vanish on all of them.

⁶²It follows directly from the way $\bar{\nabla}_a$ was defined that the pair $(\bar{\nabla}_a, \bar{\phi})$ satisfies conditions (R1) and (R2). (The argument is almost exactly the same as that used in an earlier note to prove (G1) and (G2) in the Geometrization Theorem.) What must be shown is that $\bar{\nabla}_a$ is flat, and that the pair $(\bar{\nabla}_a, \bar{\phi})$ satisfies Poisson's equation: $\bar{\nabla}^a \bar{\nabla}_a \bar{\phi} = 4 \pi \rho$. We do so by showing that (i) $\bar{R}^a_{bcd} = R^a_{bcd}$, (ii) $\bar{\nabla}^a \nabla_a \psi = 0$, and (iii) $\bar{\nabla}^a \bar{\nabla}_a \alpha = \nabla^a \nabla_a \alpha$, for all smooth scalar fields α on M . (It follows immediately from (ii) and (iii) that $\bar{\nabla}^a \bar{\nabla}_a \bar{\phi} = \bar{\nabla}^a \nabla_a \phi + \bar{\nabla}^a \nabla_a \psi = \nabla^a \nabla_a \phi + \nabla^a \nabla_a \psi = 4 \pi \rho$.)

We know from the uniqueness clause of (G2) in the Geometrization Theorem that the action of $\overset{g}{\nabla}_a$ with respect to ∇_a is given by the field $C^a_{bc} = -t_{bc} \nabla^a \phi$. It follows that the action of $\bar{\nabla}_a$ relative to that of ∇_a is given by $\hat{C}^a_{bc} = -\bar{C}^a_{bc} + C^a_{bc} = -t_{bc} \nabla^a (-\bar{\phi} + \phi) = t_{bc} \nabla^a \psi$. So, arguing almost exactly as we did in the proof of (G3) in the Geometrization Theorem, we have

$$(54) \quad \bar{R}^a_{bcd} = R^a_{bcd} + 2 t_{b[d} \nabla_{c]} \nabla^a \psi.$$

Now it follows from $\nabla^a \nabla^b \psi = 0$ that

$$(55) \quad \nabla_c \nabla^a \psi = t_c (\xi^n \nabla_n \nabla^a \psi),$$

where $t_{ab} = t_a t_b$, and ξ^n is any smooth future-directed unit timelike vector field on M . Hence, $t_{b[d} \nabla_{c]} \nabla^a \psi = t_b t_{[d} t_{c]} (\xi^n \nabla_n \nabla^a \psi) = 0$. This, together with (54), gives us (i). And (ii) follows directly from (55). Finally, for (iii), notice that

$$\begin{aligned}
 \bar{\nabla}^a \bar{\nabla}_a \alpha &= h^{ar} \bar{\nabla}_r \bar{\nabla}_a \alpha = h^{ar} \bar{\nabla}_r \nabla_a \alpha = h^{ar} (\nabla_r \nabla_a \alpha + \hat{C}^n_{ra} \nabla_n \alpha) \\
 &= \nabla^a \nabla_a \alpha + h^{ar} t_{ra} (\nabla^n \psi) (\nabla_n \alpha) = \nabla^a \nabla_a \alpha.
 \end{aligned}$$

The final equality follows from the orthogonality condition.

$-m \nabla^a \phi$. Relative to the second, it has acceleration $\xi^n \bar{\nabla}_r \xi^a = \xi^n \nabla_n \xi^a - \nabla^a \psi$ and is subject to a gravitational force $-m \nabla^a \bar{\phi} = -m \nabla^a \phi - m \nabla^a \psi$.

As suggested at the beginning of the section, we can take this non uniqueness of recovery result to capture in precise mathematical language the standard claim that Newtonian gravitational force is a gauge quantity. By the argument just given, if we can take the force on a point particle with mass m to be $-m \nabla^a \phi$, we can equally well take it to be $-m \nabla^a (\phi + \psi)$, where ψ is any field satisfying $\nabla^a \nabla^b \psi = 0$.

3.3 Recovering Global Geometric Structure from “Causal Structure”

There are many interesting and important issues concerning the global structure of relativistic spacetimes that might be considered here — the nature and significance of singularities, the cosmic censorship hypothesis, the possibility of “time travel”, and others.⁶³ But we limit ourselves to a few remarks about one rather special topic.

In our discussion of relativistic spacetime structure, we started with geometric models (M, g_{ab}) exhibiting several levels of geometric structure, and used the latter to define the (two-place) relations \ll and $<$ on M .⁶⁴ The latter are naturally construed as relations of “causal connectibility (or accessibility)”. The question arises whether it is possible to work backwards, i.e. *start* with the pair (M, \ll) or $(M, <)$, with M now construed as a bare point set, and recover the geometric structure with which we began. The question is suggested by long standing interest on the part of some philosophers in “causal theories” of time or spacetime. It also figures centrally in a certain approach to quantum gravity developed by Rafael Sorkin and co-workers. (See, e.g. Sorkin [1995; 2005].)

Here is one way to make the question precise. (For convenience, we work with the relation \ll .)

Let (M, g_{ab}) and (\bar{M}, \bar{g}_{ab}) be (temporally oriented) relativistic spacetimes. We say that a bijection $\varphi: M \rightarrow \bar{M}$ between their underlying point sets is a *causal isomorphism* if, for all p and q in M ,

$$(56) \quad p \ll q \iff \varphi(p) \ll \varphi(q).$$

Now we ask: Does a causal isomorphism have to be a homeomorphism? a diffeomorphism? a conformal isometry?⁶⁵

Without further restrictions on (M, g_{ab}) and (\bar{M}, \bar{g}_{ab}) , the answer is certainly ‘no’ to all three questions. Unless the “causal structure” (i.e. the structure determined by \ll) of a spacetime is reasonably well behaved, it provides no useful information at all. For example, let us say that a spacetime is *causally degenerate* if $p \ll q$ for all points p and q . Any bijection between two causally degenerate spacetimes qualifies as a causal isomorphism. But we can certainly find causally degenerate spacetimes whose underlying

⁶³Earman [1995] offers a comprehensive review of many of them. On the topic of singularities, I can also recommend Curiel [1999].

⁶⁴Recall that $p \ll q$ holds if there is a future-directed timelike curve that runs from p to q ; and $p < q$ holds if there is a future-directed causal curve that runs from p to q .

⁶⁵We know in advance that a causal isomorphism need not be a (full) isometry because conformally equivalent metrics g_{ab} and $\Omega^2 g_{ab}$ on a manifold M determine the same relation \ll . The best one can ask for is that it be a conformal isometry, i.e. that it be a diffeomorphism that preserves the metric up to a conformal factor.

manifolds have different topologies (e.g. Gödel spacetime and a rolled-up version of Minkowski spacetime).

There is a hierarchy of “causality conditions” that is relevant here. (See, e.g. Hawking and Ellis [1972, section 6.4].) They impose, with varying degrees of stringency, the requirement that there exist no closed, or “almost closed”, timelike curves. Here are three.

chronology: There do not exist closed timelike curves. (Equivalently, for all p , it is *not* the case that $p \ll p$.)

future (resp. past) distinguishability: For all points p , and all sufficiently small open sets O containing p , no future directed (resp. past directed) timelike curve that starts at p , and leaves O , ever returns to O .

strong causality: For all points p , and all sufficiently small open sets O containing p , no future directed timelike curve that starts in O , and leaves O , ever returns to O .

It is clear that strong causality implies both future distinguishability and past distinguishability, and that each of the distinguishability conditions (alone) implies chronology. Standard examples (Hawking and Ellis [1972]) establish that the converse implications do not hold, and that neither distinguishability condition implies the other.

The names “future distinguishability” and “past distinguishability” are easily explained. For any p , let $I^+(p)$ be the set $\{q : p \ll q\}$ and let $I^-(p)$ be the set $\{q : q \ll p\}$. Then future distinguishability is equivalent to the requirement that, for all p and q ,

$$I^+(p) = I^+(q) \Rightarrow p = q.$$

And the counterpart requirement with I^+ replaced by I^- is equivalent to past distinguishability.

We mention all this because it turns out that one gets a positive answer to all three questions above if one restricts attention to spacetimes that are *both* future and past distinguishing.

PROPOSITION 15. *Let (M, g_{ab}) and $(\overline{M}, \overline{g}_{ab})$ be (temporally oriented) relativistic spacetimes that are past and future distinguishing, and let $\varphi : M \rightarrow \overline{M}$ be a causal isomorphism. Then φ is a diffeomorphism and preserves g_{ab} up to a conformal factor; i.e. $\varphi_* g_{ab}$ is conformally equivalent to \overline{g}_{ab} .*

A proof is given in Malament [1977]. A counterexample given there also shows that the proposition fails if the initial restriction on causal structure is weakened to past distinguishability or to future distinguishability alone.

ACKNOWLEDGEMENTS

I am grateful to Jeremy Butterfield, Erik Curiel, John Earman, and Chris Smeenk for comments on earlier drafts.

BIBLIOGRAPHY

- [Bain, 2004] J. Bain. Theories of Newtonian gravity and empirical distinguishability. *Studies in the History and Philosophy of Modern Physics*, 35:345–376, 2004.
- [Budden, 1998] T. Budden. Geometric simultaneity and the continuity of special relativity. *Foundations of Physics Letters*, 11:343–357, 1998.
- [Cartan, 1923] E. Cartan. Sur les variétés à connexion affine et la théorie de la relativité généralisée. *Annales Scientifiques de l'École Normale Supérieure*, 40:325–412, 1923.
- [Cartan, 1924] E. Cartan. Sur les variétés à connexion affine et la théorie de la relativité généralisée. *Annales Scientifiques de l'École Normale Supérieure*, 41:1–25, 1924.
- [Coleman and Korté, 2001] R. Coleman and H. Korté. Hermann Weyl: Mathematician, physicist, philosopher. In E. Scholz, editor, *Hermann Weyl's Raum-Zeit-Materie and a General Introduction to his Scientific Work*. Birkhäuser Verlag, 2001.
- [Curiel, 1999] E. Curiel. The analysis of singular spacetimes. *Philosophy of Science*, 66:S119–S145, 1999.
- [Earman, 1995] J. Earman. *Bangs, Crunches, Whimpers, and Shrieks*. Oxford University Press, 1995.
- [Earman, 2001] J. Earman. Lambda: The constant that refuses to die. *Archives for History of Exact Sciences*, 55:189–220, 2001.
- [Earman, 2003] J. Earman. The cosmological constant, the fate of the universe, unimodular gravity, and all that. *Studies in the History and Philosophy of Modern Physics*, 34:559–577, 2003.
- [Ehlers and Geroch, 2004] J. Ehlers and R. Geroch. Equation of motion of small bodies in relativity. *Annals of Physics*, 309:232–236, 2004.
- [Ehlers, 1981] J. Ehlers. Über den Newtonschen Grenzwert der Einsteinschen Gravitationstheorie. In J. Nitsch, J. Pfarr, and E. W. Stachow, editors, *Grundlagen Probleme der Modernen Physik*. Wissenschaftsverlag, 1981.
- [Friedrichs, 1927] K. Friedrichs. Eine Invariante Formulierung des Newtonschen Gravitationsgesetzes und der Grenzüberganges vom Einsteinschen zum Newtonschen Gesetz. *Mathematische Annalen*, 98:566–575, 1927.
- [Geroch and Jang, 1975] R. Geroch and P. S. Jang. Motion of a body in general relativity. *Journal of Mathematical Physics*, 16:65–67, 1975.
- [Geroch, 1971] R. Geroch. A method for generating solutions of Einstein's equation. *Journal of Mathematical Physics*, 12:918–924, 1971.
- [Geroch, unpublished] R. Geroch. General Relativity. Unpublished Lecture Notes from 1972, University of Chicago, unpublished.
- [Grünbaum, 1973] A. Grünbaum. *Philosophical Problems of Space and Time*. Reidel, 2nd enlarged edition, 1973.
- [Grünbaum, forthcoming] A. Grünbaum. David Malament and the conventionality of simultaneity: A reply. In *Collected Papers*, volume 1. Oxford University Press, forthcoming.
- [Hawking and Ellis, 1972] S.W. Hawking and G. F. R. Ellis. *The Large Scale Structure of Space-Time*. Cambridge University Press, 1972.
- [Hogarth, 1993] M. Hogarth. Standard simultaneity is unique in metrical Minkowski spacetime: A generalization of Malament's theorem. unpublished, 1993.
- [Janis, 2002] A. Janis. Conventionality of simultaneity. In E. Zalta, editor, *Stanford Encyclopedia of Philosophy*, 2002. URL=<http://plato.stanford.edu/archives/fall2002/entries/spacetime-convensimul/>.
- [Künzle, 1972] H. Künzle. Galilei and Lorentz structures on space-time: Comparison of the corresponding geometry and physics. *Annales Institute Henri Poincaré*, 17:337–362, 1972.
- [Künzle, 1976] H. Künzle. Covariant Newtonian limits of Lorentz space-times. *General Relativity and Gravitation*, 7:445–457, 1976.
- [Laugwitz, 1965] D. Laugwitz. *Differential and Riemannian Geometry*. Academic Press, 1965.
- [Malament, 1977] D. Malament. The class of continuous timelike curves determines the topology of spacetime. *Journal of Mathematical Physics*, 18:1399–1404, 1977.
- [Malament, 1995] D. Malament. Is Newtonian cosmology really inconsistent? *Philosophy of Science*, 62:489–510, 1995.
- [Malament, unpublished] D. Malament. Lecture Notes on Differential Geometry. available at <http://www.lps.uci.edu/home/fac-staff/faculty/malament/FndsofGR.html>, unpublished.
- [Norton, 1995] J. Norton. The force of Newtonian cosmology: Acceleration is relative. *Philosophy of Science*, 62:511–522, 1995.
- [Norton, 1999] J. Norton. The cosmological woes of Newtonian gravitation theory. In H. Goenner, J. Renn, J. Ritter, and T. Sauer, editors, *The Expanding Worlds of General Relativity*. Birkhäuser, 1999.
- [O'Neill, 1983] B. O'Neill. *Semi-Riemannian Geometry*. Academic Press, 1983.

- [Page, 1983] D. Page. How big is the universe today? *General Relativity and Gravitation*, 15:181–185, 1983.
- [Reichenbach, 1958] H. Reichenbach. *The Philosophy of Space and Time*. Dover, 1958.
- [Rindler, 1981] W. Rindler. Public and private space curvature in Robertson-Walker universes. *General Relativity and Gravitation*, 13:457–461, 1981.
- [Sachs and Wu, 1977a] R. K. Sachs and H. Wu. General relativity and cosmology. *Bulletin of the American Mathematical Society*, 83:1101–1164, 1977.
- [Sachs and Wu, 1977b] R. K. Sachs and H. Wu. *General Relativity for Mathematicians*. Springer Verlag, 1977.
- [Sorkin, 1995] R. D. Sorkin. A specimen of theory construction from quantum gravity. In J. Leplin, editor, *The Creation of Ideas in Physics*. Kluwer, 1995.
- [Sorkin, 2005] R. D. Sorkin. Causal sets: Discrete gravity. In A. Gomberoff and D. Marolf, editors, *Lectures on Quantum Gravity (Proceedings of the Valdivia Summer School)*. Plenum, 2005.
- [Spirtes, 1981] P. Spirtes. Conventionalism in the philosophy of Henri Poincaré. unpublished Ph. D. thesis, University of Pittsburgh, 1981.
- [Stachel, 1980] J. Stachel. Einstein and the rigidly rotating disk. In A. Held, editor, *General Relativity and Gravitation. One Hundred Years After the Birth of Albert Einstein*, volume I. Plenum, 1980.
- [Stein, 1991] H. Stein. On relativity theory and the openness of the future. *Philosophy of Science*, 58:147–167, 1991.
- [Trautman, 1965] A. Trautman. Foundations and current problems of general relativity. In S. Deser and K. W. Ford, editors, *Lectures on General Relativity*. Prentice Hall, 1965.
- [Wald, 1984] R. Wald. *General Relativity*. University of Chicago Press, 1984.
- [Weyl, 1921] H. Weyl. Zur Infinitesimalgeometrie: Einordnung der projektiven und konformen Auffassung. *Nachrichten der Königlichen Gesellschaft der Wissenschaften zu Göttingen; Mathematisch-Physikalische Klasse*, pages 99–112, 1921.
- [Weyl, 1950] H. Weyl. *Philosophy of Mathematics and Natural Science*. Princeton University Press, 2nd edition, 1950.

This page intentionally left blank

NON-RELATIVISTIC QUANTUM MECHANICS

Michael Dickson

This article is an introduction to some of the most important philosophical and foundational issues that arise from or concern non-relativistic quantum theory. The chapter has six main sections. The first introduces the theory, including some of the important mathematical results required to formulate and address many of the philosophical and foundational issues. This section is the longest, and most important, for it will begin to give the careful reader the background needed to understand and evaluate much of the vast literature on non-relativistic quantum theory. And that literature is indeed vast — there is no way that it can even be summarized in a chapter of this length. Instead, in the five subsequent sections, I will consider some of the more important issues: foundational characterizations of the formalism of quantum theory, empirical content, quantum uncertainty, the measurement problem, and non-locality. There are many other issues one could discuss, and some recent movements that merit consideration. Alas, we will not have time for them. A careful reading of the material here is a start, however, towards understanding these other issues.

Of these five issues, the first two are somewhat less discussed, especially in the Anglo-American philosophical literature. Those sections are therefore longer, relative to the final three, than some readers might expect. This fact is not meant to imply anything about the relative importance of the issues, but is an attempt to redress a relative lack of coverage in certain circles.

Much of the material presented here — especially from §4, §5, and §6 — is largely my review of standard material that can be found in many places. I have therefore chosen not to provide extensive bibliographic information. Indeed, I have kept bibliographic references to a minimum. This article is thus *not* intended to be a compendium of work in the field, much less an extensive annotated bibliography. The reader is encouraged to seek additional resources to fill out the brief accounts given here. Such resources are numerous.

The final section is a brief mathematical appendix, reviewing essential definitions and results, mostly from the theory of Hilbert spaces and groups. It may serve one of two purposes, depending on the reader: a brief reminder of concepts learned elsewhere; or a prompt to learn the concepts elsewhere. It is unlikely that a reader who is completely unfamiliar with these concepts will absorb them just from what is said here. Reference is made to the relevant subsections of this appendix at the appropriate places in the text.

While at some points I have made some effort at rigor, for the most part, the discussion here is only partly rigorous, with the occasional attempt made to point towards what further would be required for complete rigor. The reader is, again, encouraged to consult the literature for mathematical details, and in any case is encouraged to bear in mind that much of the discussion here is *not* intended to be entirely mathematically rigorous, while, I hope, also not being misleading.

1 THE THEORY

This section is an introduction to the formalism of quantum theory. After a brief justification of this approach (§1.1), I will introduce the major elements of the formalism (§1.2), followed by a simple, but important, example (§1.3). I will then introduce the commonly used ‘Dirac’ notation (§1.4), and conclude by considering the role of transformations (groups) in the theory (§1.5), including dynamical transformations (equations of motion) and finally (§1.6) a brief preview of the philosophical issues to come. More than subsequent sections, this section will rely heavily on the material from the mathematical appendix (§7), with references where appropriate.

1.1 *The Thought Behind Starting with Formalism*

Why begin an account of a physical theory with its formalism? Why not begin, instead, with its basic physical insights, or fundamental physical principles? One problem with such an approach here is that, in the case of quantum theory, there is not much significant agreement about what the basic physical insights, or fundamental physical principles, are. Some argue that the collapse postulate (to be discussed later) is at the heart of the theory. Others argue that it must be excised from the theory. Some argue that the theory is fundamentally indeterministic, while others argue that we can make sense of it only in terms of an underlying determinism. Some argue that the familiar notion of a ‘particle’ with a definite location is a casualty of the theory, while others argue that the theory makes sense only if one takes such a notion as fundamental.

Now, advocates of these different views tend also, it is true, to advocate different formulations of the theory, but they will not suggest that formulations other than their preferred one are *wrong*, only that they, perhaps, emphasize the wrong points. (Indeed, there is no disputing that the standard formalism — the one presented here — is empirically successful; advocates of different views will ultimately have to account for that success in their own terms.) Hence, while the choice of a single formalism at the start of our discussion might slant our point of view somewhat, it will, unlike the choice of basic physical insights or fundamental physical principles, not prejudice the central issues.

1.2 The Standard Formalism

I begin with a very brief sketch of a common understanding of the formalism, which I shall flesh out and generalize subsequently. (The reader is not expected to have a deep understanding of any aspect of the formalism merely as a result of reading this subsection.)

1.2.1 Hilbert Space

The formalism of quantum mechanics is normally understood in terms of the theory of Hilbert spaces (§7.1). A Hilbert space is a vector space (§7.1.1) with an inner product (§7.1.3) that is also complete with respect to the norm (§7.1.4) defined by this inner product. A standard example is the space, ℓ^2 , of (modulus)-square-summable sequences of complex numbers. In this space, the inner product of two vectors, (x_1, x_2, \dots) and (y_1, y_2, \dots) is $\sum_{n=1}^{\infty} x_n^* y_n$. Another standard example is the space, $L^2(\mathbb{R}^N)$, of (modulus)-square-integrable, Lebesgue-measurable, complex-valued functions on \mathbb{R}^N , where we identify two functions (i.e., they represent the same vector) if (and only if) they differ only on a set of Lebesgue measure (§7.5.4) zero. Here the inner product of two vectors, $f(x)$ and $g(x)$, is $\int f^*(x)g(x)dx$ (where $f(x)$ and $g(x)$ are arbitrary representatives from their respective equivalence classes).¹

1.2.2 Observables

The ‘observables’ of the theory — the physical quantities, or properties, whose value or presence one can, in principle at least, measure, or ‘observe’ — are normally taken to be represented by the self-adjoint operators (§7.2.1, §7.2.3) on the Hilbert space. (The nature of the representation — that is, which operators represent which observables — can depend on the physical situation being described.) Via the spectral theorem (discussed below), one can identify each observable with a spectral family of projection operators, the observable being given, essentially, by a map from Borel sets (§7.5.5) of possible values of the observable to elements in the spectral family. This subsection reviews these ideas briefly.

1.2.2.1 Positive Operator Valued Measures It is often useful to adopt a broader notion of an observable, as a ‘positive-operator-valued measure’ (POVM). In this approach, we begin with a set of ‘possible values’ for the observable, represented in the most general case as a locally compact topological space, S (§7.5.1). In most cases of interest to us, S is a subset of the real numbers, or things can be reworked so that it is.

¹For those who have some familiarity with quantum theory: the space ℓ^2 is the space used in Heisenberg’s ‘matrix mechanics’, while the space $L^2(\mathbb{R}^N)$ is the space used in Schrödinger’s ‘wave mechanics’. As Hilbert spaces, ℓ^2 and $L^2(\mathbb{R}^N)$ are isomorphic, meaning that the two theories are essentially the same.

A map, $E : \mathcal{B}(S) \rightarrow \mathcal{B}(\mathcal{H})$, from the Borel subsets of S to the bounded operators (§7.2.2) on some Hilbert space, \mathcal{H} , is a POVM just in case for any disjoint sequence of such subsets, $\Delta_n \subseteq S$,

- (1) $E(\Delta_n)$ is a positive operator for all n
- (2) $E(S) = \mathbb{I}$, the identity on \mathcal{H}
- (3) $E(\cup_n \Delta_n) = \sum_n E(\Delta_n)$.

In (1), an operator, E , is positive if $\langle v, Ev \rangle \geq 0$ for all $v \in \mathcal{H}$. The positive operators on \mathcal{H} are denoted by $\mathcal{B}(\mathcal{H})^+$. The convergence intended in (3) is in the weak operator topology on \mathcal{H} (§7.5.3). If, in addition, $E(\Delta_n \cap \Delta_m) = E(\Delta_n)E(\Delta_m)$ whenever $n \neq m$ then: everything in the image of E is a projection operator; E is then called a ‘projection-valued measure’ (PVM); and the family $\{E(\Delta_n)\}$ is a ‘spectral family’. In this case, the $E(\Delta_n)$ are mutually orthogonal, meaning that $E(\Delta_m)E(\Delta_n) = 0$ (the zero operator) whenever $m \neq n$, and we write $E(\Delta_m) \perp E(\Delta_n)$.

We can recover a self-adjoint operator from any PVM, E . If the cardinality of $S \subset \mathbb{R}$ is finite ($S = \{s_1, \dots, s_N\}$), then the recovery is straightforward:²

$$(4) \quad F = \sum_{n=1}^N s_n E(s_n).$$

That is, the operator F is the weighted sum of the (mutually orthogonal) projections $E(s_n)$, the weights being the ‘possible values’ of the observable, i.e., elements of S . If S is countably infinite, then the situation is much the same, though one must worry about convergence. If S is uncountably infinite, then the sum becomes an integral, and matters become considerably more complicated. In any case, the resulting operator, F , is self-adjoint.

1.2.2.2 Spectral Theorem The spectral theorem states the converse of the construction given by (4). Again, the finite case is simplest. There, every self-adjoint operator, F , can be written as

$$(5) \quad F = \sum_n s_n P_n,$$

where the s_n are real numbers and the P_n are mutually orthogonal projections. A spectral family therefore fixes a self-adjoint operator, and a self-adjoint operator fixes a spectral family. Hence the formalism of PVMs makes quick contact with a formalism (in terms of self-adjoint operators) that is perhaps more familiar to

²In (4) $E(s_n)$ should strictly be written $E(\{s_n\})$, because $E()$ acts on Borel sets, but the notation is clumsy, and the meaning of $E(s_n)$ should be clear enough. I follow the same convention elsewhere.

some readers, and certainly widely used in physics; thus that latter formalism can be seen as a special case of the more general formalism in terms of POVMs. The case of infinite-dimensional spaces is conceptually analogous, but mathematically trickier.

Notice that every vector inside the subspace corresponding to a projection P_n (henceforth, $\text{ran } P_n$, the ‘range’ of P_n) is an eigenvector (§7.2.1) of F ; the $\text{ran } P_n$ are therefore often called ‘eigenspaces’ of F . When the eigenspaces of F are all one-dimensional, then F is called ‘maximal’. The import of being maximal will become clear below.

Finally, notice that the spectral projections of F partially define an orthogonal, indeed orthonormal (§7.1.4), basis for the space. Within each P_n , choose a set of mutually orthogonal and normalized vectors, $\{e_{n,m}\}_{m=1}^{\dim(P_n)}$. Do the same for the kernel of F (the subspace that F maps to 0, denoted $\ker F$). The result (i.e., the union of all these sets) is an orthonormal basis (which, if F is maximal, is in fact fixed up to constant multipliers of the elements of the basis). Even when this basis is not uniquely fixed by F (because it is non-maximal), I will refer to such a basis as ‘a basis determined by F ’.

1.2.3 States

1.2.3.1 Probabilities The formalism in terms of POVMs (as well as the special case of PVMs) describes a probabilistic theory, inasmuch as it provides probabilities for (Borel sets of) values of observables, or (equivalently and sometimes more conveniently) expectation values for observables. I will take probabilities as fundamental; expectation values can then be generated from a probability measure over the possible values, f_n , of F in the usual way:

$$(6) \quad \text{Exp}(F) = f_1 \text{Pr}(f_1) + f_2 \text{Pr}(f_2) + \dots$$

As we noticed above, rather than considering directly the possible values of an observable, we can also consider the corresponding (spectral) projections, which can be taken, in a given physical situation, to represent those values.

A probability measure, p , defined on the projection operators should, minimally, be such that $p(P_1 + P_2) = p(P_1) + p(P_2)$ whenever $P_1 \perp P_2$. (Later I will motivate this condition. The basic idea is that it corresponds to the usual ‘additivity axiom’ of Kolmogorovian probability theory — see §7.5.6.) More specifically, and for now considering just the case of PVMs, we require a probability measure on the projections on a Hilbert space to be a map, p , from projections to the interval $[0, 1]$, where p is countably additive on sets of mutually orthogonal projections.

Precisely what one means by countable additivity for the operators that are in the image of a POVM (rather than a PVM) is a slightly subtle matter. In particular, in general the operators in the image of a POVM — normally they are called ‘effects’ — do not correspond to subspaces, and the notion of orthogonality does not apply. However, there is a natural generalization of the concept. Notice that for projections, $\{P_i\}$, in the image of a PVM, the condition that $\mathbb{I} - \sum_i P_i$

be a projection (or maybe the zero operator) is equivalent to the condition that the $\{P_i\}$ be mutually orthogonal.³ The analogous condition in the case of positive operators is that, for effects $\{E_i\}$ in the image of a POVM, if $\mathbb{I} - \sum_i E_i$ is positive (or 0), then $\Pr(\sum_i E_i) = \sum_i \Pr(E_i)$.

1.2.3.2 Statevectors and Wavefunctions

1.2.3.2.a Statevectors Normalized vectors determine probability measures over the projections, via:

$$(7) \quad \text{probability of } P \text{ as given by } v := \Pr_v(P) := \langle v, Pv \rangle.$$

(One often sees the expression $|\langle \psi, v \rangle|^2$, where ψ is a normalized vector from $\text{ran } P$. The two expressions are equivalent.) Notice that the probabilities generated by the vectors v and $e^{i\phi}v$ (where ϕ is a real number) are the same. One says that ‘overall phases do not affect probabilities’. The expectation value of a self-adjoint operator, F , given by the state v is

$$(8) \quad \text{expectation of } P \text{ as given by } v := \text{Exp}_v(F) := \langle v, Fv \rangle.$$

(Note that the expectation value of a projection is also its probability.)

Note that if $v \in \text{ran } P$ then $\Pr_v(P) = 1$. More generally, if P_n is an eigenspace of F corresponding to the eigenvalue f_n and $v \in \text{ran } P_n$, then $\Pr_v(P_n) = 1$, i.e. the probability (in the state v) that F has the value f_n is 1. Such a state, v , is called an ‘eigenstate’ of F — it is a normalized vector inside the eigenspace, $\text{ran } P_n$, of F . Notice that, in this case, writing F in terms of its spectral decomposition (recall 4) makes the determination of probabilities and expectation values trivial. Indeed, even when dealing with general states, it is often convenient to write F in terms of its spectral decomposition, and the state in terms of a basis determined by F .

1.2.3.2.b Superposition It is a standard assumption of quantum theory that *every* vector in the Hilbert space for a system is a possible state for the system. This assumption is often expressed as the ‘superposition principle’, which asserts that (normalized) linear combinations of statevectors are again statevectors.

Given an observable, F , the superposition principle gives rise to (possible) states that are *not* eigenstates of F . Suppose, for simplicity, that F is maximal, with eigenspaces and eigenvalues $\{P_n\}$ and $\{f_n\}$, and consider an orthonormal basis, $\{v_n\}$, determined by F (which, because F is maximal, just amounts to choosing one normalized vector from each $\text{ran } P_n$). Now form the statevector

$$(9) \quad v = \sum_n k_n |v_n\rangle$$

³Sketch of a proof: Write $(\mathbb{I} - \sum_i P_i)(\mathbb{I} - \sum_i P_i)$; expand; argue that if the $\{P_i\}$ are mutually orthogonal, then the result is $\mathbb{I} - \sum_i P_i$; argue (using the fact that projections are positive — this part is less trivial) that if the result is $\mathbb{I} - \sum_i P_i$, then the $\{P_i\}$ are mutually orthogonal; finally, argue that $\mathbb{I} - \sum_i P_i$ is self-adjoint.

where $\sum_n |k_n|^2 = 1$ and with at least *two* non-zero coefficients k_n . In this case, we say that v is a *superposition* of the v_n . (One sometimes hears the word ‘superposition’ used in a way that suggests that some vectors are ‘in superpositions’ and others are not. *Relative to a given basis*, this distinction makes sense, but otherwise it does not. *Every* vector is a superposition for some choices of basis.) Notice that v is not an eigenstate of F , and assigns non-trivial probabilities to more than one possible value of F . Of course, the superposition principle implies that v is nonetheless a possible state of a system.

1.2.3.2.c Wavefunctions Wavefunctions are just a specific way of representing statevectors. It is often convenient to take the Hilbert space for a quantum system to be the elements of $L^2(\mathbb{R}^3)$, in which case statevectors are (equivalence classes of) complex-valued functions on \mathbb{R}^3 . The equation of motion that they standardly satisfy is a type of wave equation (e.g., the Schrödinger equation — see §1.5.2.3.a), and for this reason — as well as the fact that the equation was historically derived with wave phenomena in mind — these functions are called ‘wavefunctions’. Linear combinations of waves may be conceived in terms of ‘superposing’ the waves; hence the term ‘superposition’.

1.2.3.3 Gleason’s Theorem One can generate probability measures using non-negative trace-1 operators (‘density operators’). The functional $\text{Tr}[\cdot]$ is the ‘trace functional’, a map from the bounded operators on a Hilbert space to \mathbb{R} defined by:

$$(10) \quad \text{Tr}[F] = \sum_k \langle e_k, (F^*F)^{1/2} e_k \rangle$$

where $\{e_k\}$ is an orthonormal basis for \mathcal{H} . (Note that F^*F is self-adjoint and positive. It is in fact true that every positive operator, A , has a positive self-adjoint square root, B , defined by $B^2 = A$.) And if F itself is positive, then $F = \sqrt{F^2}$ and

$$(11) \quad \text{Tr}[F] = \sum_k \langle e_k, F e_k \rangle.$$

The trace functional is provably independent of the choice of orthonormal basis, $\{e_i\}$. Moreover, a very useful property of the trace functional is that it is invariant under cyclic permutations of its arguments; for example,

$$(12) \quad \text{Tr}[ABC] = \text{Tr}[BCA] = \text{Tr}[CAB]$$

for any A, B, C .

Let W be any positive operator on a Hilbert space, \mathcal{H} , with $\text{Tr}[W] = 1$. Let $E(\cdot)$ be any POVM from some ‘spectrum’, S , of possible values to positive operators. Then $\text{Tr}[WE(\cdot)]$ is a countably additive probability measure on (the σ -algebra of

Borel sets of) possible values of the observable represented by the POVM E as follows:

$$(13) \quad \Pr(\Delta) = \text{Tr}[WE(\Delta)].$$

Countable additivity follows from (3) and the linearity of the trace functional. Normalization follows from (2) and the fact that W has unit trace.

When $E(\cdot)$ is a PVM, (13) defines a countably additive normalized measure on the projections on \mathcal{H} . Hence any density operator generates such a measure. The converse is (remarkably) true as well: every probability (i.e., countably additive, normalized) measure on the projections on a Hilbert space is generated as in (13) by some density operator. This theorem is due to Gleason ([1957]), and says, more precisely:

Theorem (Gleason): Let \mathcal{H} be a Hilbert space of dimension greater than 2. Then every countably additive normalized measure, $\Pr(\cdot)$, on the projections on (equivalently, closed subspaces of) \mathcal{H} is generated by some trace-1 positive operator, W , on \mathcal{H} ; for P a projection,

$$(14) \quad \Pr(P) = \text{Tr}[WP].$$

The proof is non-trivial. Gleason's theorem is generalizable to the case of general POVMs. That is, the countably additive probability measures over effects are also given by the density operators. (Indeed, for POVMs, there is no restriction to the case $\dim(\mathcal{H}) > 2$. Again, the proofs are non-trivial. See Busch ([2003]).)

In this common understanding of quantum theory, then, the kinematics of a quantum system is, at its core, given by the POVMs on a Hilbert space together with a state, a density operator. In many cases of interest, one deals with PVMs, hence self-adjoint operators, rather than with POVMs.

Note, finally, that for any statevector, v , we can always represent v in terms of the density-operator formalism, by choosing as the state the projection, P_v , onto the subspace spanned by v . In this case, for any projection Q , $\text{Tr}[P_v Q] = \langle v, Qv \rangle$. (To prove: take the trace in an orthonormal basis containing v .)

1.2.3.4 Matrix Representation of States A vector — and in particular a statevector, ψ — can, of course, be written in terms of any orthonormal basis, $\{e_n\}$, and in this case, the coefficients c_n in the expansion $\psi = \sum_n c_n e_n$ may be considered as the ‘coordinates of ψ in the e_n -basis’. It is, in fact, sometimes convenient (see, e.g., §1.3.3.2) to write the state as a column vector with these coordinates.

A similar construction is available for density operators. Again in the (orthonormal) basis $\{e_n\}$, consider a matrix whose elements are $\langle e_n, F e_m \rangle$, for any operator, F , on a Hilbert space, \mathcal{H} . This map from operators on \mathcal{H} to $N \times N$ matrices (where N could be infinite) is in fact an isomorphism from the (algebra of) operators on \mathcal{H} to the (algebra of) $N \times N$ matrices

In particular, let W be a density operator on \mathcal{H} , and let $W_{nm} = \langle e_n, F e_m \rangle$. Now let F be an observable whose eigenvectors are the e_n . Notice, in this case,

that $\langle e_n, F e_m \rangle = \delta_{nm}$. One says that F is ‘diagonal’ in the basis $\{e_n\}$ (because all of the entries off of the diagonal are 0). If W is *also* diagonal in $\{e_n\}$, then the probabilities assigned by W to F behave completely classically, and in particular the classical ‘sum rule’ holds:

$$(15) \quad \Pr_W(f_n \text{ or } f_m) = \Pr_W(f_n) + \Pr_W(f_m)$$

(where \Pr_W is the probability assigned by W via 14 and f_n is the eigenvalue of F corresponding to the eigenvector e_n). However, if W is *not* diagonal in $\{e_n\}$, then in general (15) fails. In this case, one speaks of ‘interference’ between the e_n (in the state W).

1.2.3.5 Expectation Values It follows immediately that the expectation value of the observable (represented by the self-adjoint operator) F in the state (represented by the density operator) W is $\text{Tr}[WF]$. To see why, write F in terms of its spectral resolution. The point is most easily seen when F has only a discrete spectrum, as in (5). Then by the linearity of the trace,

$$(16) \quad \text{Tr}[WF] = \sum_n \text{Tr}[WP_n]s_n.$$

(When F has a continuous spectrum, one must work with integrals whose definition must be treated carefully.) Notice that the expression $\text{Tr}[WP_n]$ is the probability (in state W) that F takes the value s_n . Hence (16) is a weighted sum of the possible (spectral) values, s_n , for F , the weights given by the probabilities, $\text{Tr}[WP_n]$, associated to those values in the state W . Note that the traces in (16) will in general be easiest to calculate in a basis determined by F .

1.2.3.6 Quantum Probability Theory Classical probability theory standardly concerns measures over sigma-algebras of events (§7.5.5, §7.5.6). These sigma-algebras are defined in terms of the usual set-theoretic operations of complement and union. In quantum theory, we are dealing with a different structure. However it is sufficiently analogous to the structure considered in the classical setting that, mathematically at least, one can often easily carry over considerations from classical probability theory. Our ‘sample space’ is the set of all one-dimensional projections. Set-theoretic complement (E') becomes ‘orthogonal complement’ (E^\perp); set-theoretic union ($E \cup F$) becomes ‘span’ (the span of the subspaces E and F , written $E \vee F$); set-theoretic intersection ($E \cap F$) remains intersection (now written $E \wedge F$); and set-theoretic ‘inclusion’ ($E \subseteq F$) becomes subspace inclusion (often written $E \leq F$). Later, I will consider this structure in more detail — it is the ‘lattice’, \mathcal{L} , of subspaces of a Hilbert space (§7.4). For now, I simply note that it has the correct properties: (i) $\mathcal{H} \in \mathcal{L}$, (ii) $E \in \mathcal{L}$ implies $E^\perp \in \mathcal{L}$; and (iii) for any countable sequence, $\{E_k\} \in \mathcal{L}$, $\vee_k E_k \in \mathcal{L}$. Analogous to classical probability theory, quantum probability theory is then the theory of normalized measures on such a structure. (Of course, if we are thinking in terms of POVMs rather than PVMs, then this story cannot be told, at least not in

its present form. Instead, one considers the algebra of effects, and probability measures over it. However, I will not pursue the details here.)

1.2.3.7 Lüder's Rule What about conditional probabilities? Although its interpretation can be highly contentious, and its application somewhat tricky, there is a standard expression for a conditional probability in quantum theory, called 'Lüder's Rule'. Indeed, one can derive it from elementary considerations.

Recall from basic probability theory that the conditional probability, $\Pr(A|B)$, of one event, A , given another, B , is defined by

$$(17) \quad \Pr(A|B) := \frac{\Pr(A \cap B)}{\Pr(B)}.$$

The thought behind this definition is that the probability of A (and B) *given* B is the probability that A and B occur jointly, 'renormalized' under the assumption that B occurred; i.e., it is the probability of A 'as if' B had probability 1. Indeed, (17) is the only probability measure that satisfies the condition that if $A \subseteq B$ then $\Pr(A|B) = \Pr(A)/\Pr(B)$. In other words, if A is contained in B , then $\Pr(A|B)$ is just a renormalization of the original probability measure to one that assigns probability 1 to B .

It turns out that this condition is already sufficient to determine the form of the conditional probability measure over the (lattice of) closed subspaces of (or projections on) a Hilbert space (Bub [1977]). In other words, let \Pr_W be the probability measure associated with the density operator, W , on \mathcal{H} . Let P be a subspace such that $\Pr_W(P) \neq 0$ (where, of course, $\Pr_W(P) = \text{Tr}[WP]$). Then there is a unique probability measure, $\Pr_{W|P}$ (the 'probability in state W conditional on P '), over the closed subspaces of \mathcal{H} such that

$$(18) \quad \Pr_{W|P}(Q) := \Pr_W(Q|P) = \frac{\Pr_W(Q)}{\Pr_W(P)}$$

for any $Q \leq P$. That measure is given by

$$(19) \quad \Pr_W(Q|P) = \frac{\text{Tr}[PWPQ]}{\text{Tr}[WP]}.$$

(19) is known as 'Lüder's Rule'. Note that for a statevector, $|v\rangle$, the same effect is achieved by projecting $|v\rangle$ onto P , normalizing the result, and using that new state ($P|v\rangle/||Pv||$) to calculate the probability of Q . Hence (using eq. 7)

$$(20) \quad \Pr_{|v\rangle}(Q|P) = \langle Pv|QPv\rangle/||Pv||^2.$$

1.2.3.8 Mixed Versus Pure States Density operators that correspond to one-dimensional projections (equivalently, statevectors) are 'pure' states. These states assign probability 1 to that one-dimensional projection. Mixed states (i.e., states that are not pure) do not assign probability 1 to *any* one-dimensional projection. Moreover, mixed states are called 'mixed' because they can always be

written as a linear combination of pure states. Indeed, by the spectral theorem, any mixed state, W , can be written as $W = \sum_n w_n P_n$ (and because W is a density operator, $0 \leq w_n \leq 1$ and $\sum_n w_n = 1$). If one or more of the P_n is not one-dimensional, we can always write it as a sum of mutually orthogonal one-dimensional projections — so we may assume, without loss of generality, that all of the P_n are one-dimensional.

The coefficients, or ‘weights’, w_n , must add to one (because $\text{Tr}[W] = 1$), and in fact w_n is the probability assigned by W to P_n . Hence one can apparently think of W as representing, literally, a ‘mixture’ of systems in the pure states P_n , in the proportions w_n , so that w_n is the probability that a system chosen at random from the mixture will be found in the (pure) state P_n . We will explore (and qualify) this interpretation of mixed states below.

The converse is also true: any convex combination of pure states is again a state, in general mixed. Indeed, consider the operator

$$(21) \quad W = \sum_n w_n P_n$$

where the P_n are here one-dimensional but *not* necessarily mutually orthogonal (yet still, $\sum_n w_n = 1$).⁴ This W has unit trace (because the trace functional is linear), and therefore it is a density operator. Note, however, that (21) is in general *not* its spectral decomposition.

1.2.3.9 The Eigenstate-Eigenvalue Link According to a standard interpretation of quantum states, a system in the state W has a value for the observable F if and only if W assigns probability 1 to one of the possible values of F (and 0 to the others — in other words, ‘trivial probabilities’).⁵ Notice, in particular, that this interpretation of states differs from the usual interpretation of classical probabilistic states. In the classical case, the probabilistic state is a measure over possible pure states, and one normally presumes that the system really is in one of those pure states.

This rule for assigning definite values has come to be called, following Fine ([1973]), the ‘eigenstate-eigenvalue link’. Later (§5) we will consider in some detail the apparent consequences of this rule.

1.2.4 Incompatibility

An immediate consequence of this formalism is the fact that there are ‘incompatible’ physical quantities, at least in the minimal sense that if a state assigns

⁴More generally, if some of the P_n are not one-dimensional, then we require that $\sum_n w_n \dim P_n = 1$, because in general, for a projection P , $\text{Tr}[P] = \dim P$.

⁵For unbounded observables, such as position and momentum, one is naturally motivated to seek some other account. One possibility is to deny that they ever have definite values, but consider instead coarse-grained values, asserting, for example, that if the state W assigns probability 1 to some *region*, Δ , then the system is definitely confined to Δ , where this latter assertion is not meant to imply that there is some *point* in Δ that is the location of the system. There are, however, other approaches. See, for example, Halvorson ([2001]).

probability 1 to some physical quantity (some projection, for example), then it necessarily assigns non-trivial probabilities (i.e. neither 0 nor 1) to others (and then, by the eigenstate-eigenvalue link, these other observables do not *have* values, in that state — recall §1.2.3.9). This fact follows directly from Gleason's theorem. (Note, however, that one can show in other, simpler, ways that there are no two-valued probability measures over the projections on a Hilbert space.)

Incompatibility is closely related to non-commutativity, and indeed the two terms are sometimes used interchangeably. Consider two projection operators, Q and Q' . To keep things simple, we will suppose throughout that Q and Q' are one-dimensional. Then if Q and Q' do not commute, i.e., $[Q, Q'] \neq 0$, there is no state that assigns probability 1 to Q and either 0 or 1 to Q' . To prove this claim, we will first show (next paragraph) that the only state assigning probability 1 to a one-dimensional projection, Q , is the state Q itself. (Notice that in the previous sentence, the first mention of Q is as the representative of some physical quantity, and the second is as a state.) We will then show (subsequent paragraph) that Q assigns non-trivial probabilities to any non-commuting Q' .

Let W be a state that assigns probability 1 to (one-dimensional) Q . Writing W in terms of its spectral decomposition, and taking the trace in a basis determined by W , we immediately find that

$$(22) \quad \text{Tr}[WQ] = \sum_n w_n \langle e_n, Qe_n \rangle = 1$$

where the weights w_n (from the spectral decomposition of W) sum to 1. Hence for some n , $Qe_n = e_n$, i.e., W is in fact pure, and equal to Q . Therefore, the only state assigning probability 1 to a one-dimensional projection, Q , is Q itself.⁶

Now suppose that (one-dimensional) $Q' \neq Q$ and $Q' \not\perp Q$, i.e., Q and Q' do not commute (for a discussion, see below). Then, by the same reasoning as above, replacing Q with Q' in (22), if $\text{Tr}[WQ'] = 1$ then W must be pure and lie inside the subspace associated with Q' ; i.e., $W = Q'$. But it cannot, because we assumed that $Q \neq Q'$. On the other hand, if we want $\text{Tr}[WQ'] = 0$, then $\ker W \subseteq \text{ran } Q'$. (The reasoning is essentially the same as above.) But again it cannot, because then $Q' \perp Q$, given our earlier conclusion that W is pure and lies in the subspace associated with Q , and we already assumed that $Q' \not\perp Q$.

This fact is also true in a more general form. Given two self-adjoint operators, F and G , if F and G do not share any eigenvectors then any state that assigns probability 1 to some value for F will necessarily assign non-trivial probabilities (neither 0 nor 1) to more than one of the possible values of G . I leave the proof (using essentially the same reasoning as above) to the reader.

Above I claimed that (one-dimensional) Q and Q' do not commute if $Q' \neq Q$ and $Q' \not\perp Q$. In fact, the following is true. For any subspaces, A and B , and the

⁶This claim is also true in a more general form. Let the state W assign probability 1 to the projection Q (of any dimension). Then $(\text{ran } Q)^\perp \subseteq \ker(W)$, with equality if Q is the smallest subspace to which W assigns probability 1.

corresponding projections P_A and P_B , $[P_A, P_B] = 0$ if and only if

$$(23) \quad A = (A \wedge B) \vee (A \wedge B^\perp) \quad \text{and} \quad B = (B \wedge A) \vee (B \wedge A^\perp).$$

(Here we are not restricting to one-dimensional subspaces. Note, however, that (23) is implied by the disjunction ‘ $A = B$ or $A \perp B$ ’, and for one-dimensional subspaces, they are equivalent.) Here is the idea of the proof. Note that $A \wedge B$ and $A \wedge B^\perp$ are orthogonal. Hence, if (23) holds, we may write $P_A = P_Z + P_{A'}$ for some $Z \perp A'$. (Indeed, of course, $Z = A \wedge B$ and $A' = A \wedge B^\perp$.) Similarly, $P_B = P_Z + P_{B'}$, with $B' \perp Z$. Moreover, $A' \perp B'$. In other words, the conditions (23) imply that A and B ‘are orthogonal apart from some shared part (Z)’. Then $[P_A, P_B] = [P_Z + P_{A'}, P_Z + P_{B'}] = [P_Z, P_Z] + [P_Z, P_{B'}] + [P_{A'}, P_Z] + [P_{A'}, P_{B'}] = 0$.

Going the other way, we will just sketch the idea. If P_A and P_B commute, then for any vector, v , $P_A P_B v = P_B P_A v$. First choose $v \in A$, so that $P_A P_B v = P_B v$. In general, if $P_A w = w$ (here $w = P_B v$), then either $w \in A$ or $w = 0$. Hence either (i) $P_B v = 0$, or (ii) $P_B v \in A$. If (i) is true for all $v \in A$, then $B \perp A$ and (23) clearly holds. If (ii) holds for all $v \in A$ then $B \leq A$ and again (23) clearly holds. Using the linearity of the operators involved, one can show that if (ii) holds for just *some* $v \in A$, then the $P_B v$ must form a subspace of A , and clearly this subspace is common to A and B ; indeed it is $A \wedge B$. Similarly, one can show that choosing v from the subspace orthogonal to $A \wedge B$ gives rise to (i), so that indeed $A = (A \wedge B) \vee (A \wedge B^\perp)$. Repeating the argument for $v \in B$, we find that (23) holds.

The fact of incompatibility marks a significant departure from classical physics, where the structure of the space of states and observables allows for states that assign values to all observables with probability 1 (i.e., there are two-valued probability measures over the space of all ‘properties’ of the system). The probabilities of quantum theory appear, therefore, to be of a fundamentally different character from the probabilities of classical theory, which arise always because the state of the system is not maximally specific.⁷

1.2.5 Canonical Commutation Relations

An important and classic example of incompatibility involves the position and momentum observables. In fact, they obey the ‘canonical commutation relations’ (CCRs):

$$(24) \quad [P_i, Q_j] = -i\delta_{ij}$$

where i and j can be x , y , or z . (Henceforth, we will restrict our attention to one dimension, writing $[P, Q] = -i$. The generalization to three dimensions is straightforward.) Note that the constant on the right-hand side implicitly multiplies the identity operator.

⁷Here we are considering just cases where classical physics delivers genuine probability measures, and we ignore cases where classical physics is simply indeterminate. See Earman, Ch. 15, this volume.

Any two observables that obey these commutation relations are typically called ‘canonically conjugate’. These relations are central in quantum theory, and we will discuss them in detail in §4. For now, we simply notice them as a central example of incompatibility.

1.2.6 Compound Systems

1.2.6.1 Entangled States Compound systems are represented by tensor-product Hilbert spaces (§7.1.9), so that, for example, a system composed of two particles has a state that is a density operator on the tensor-product of the Hilbert spaces for the two particles individually. There is a fundamental and physically crucial distinction between two kinds of vector in $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$. A vector, v , in \mathcal{H} is called ‘factorizable’ if it can be written as $x \otimes y$ for some $x \in \mathcal{H}_1$ and $y \in \mathcal{H}_2$. Otherwise, v is called ‘unfactorizable’, or ‘entangled’. An analogous definition applies to the operators (hence, the density operator states) on \mathcal{H} .

The existence of entangled states (whether represented as density operators or vectors) turns out to have numerous interesting consequences. It is connected with ‘quantum non-locality’, as well as the possibility of certain computational and information-theoretic (for example, cryptographic) feats that cannot be done with classical systems.⁸ The existence of these states follows from the demand that the pure (vector) states for the compound system be closed under taking linear combinations. In other words, it follows from applying the superposition principle to compound systems as well as to simple systems.

1.2.6.2 Biorthogonal Decomposition An important result about vectors in tensor-product spaces is the ‘biorthogonal decomposition theorem’ (Schrödinger [1935b]), which states that, given a vector, v , in a Hilbert space, \mathcal{H} , and a factorization of \mathcal{H} as $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$, there exist orthonormal bases $\{e_n\}$ of \mathcal{H}_1 and $\{f_m\}$ of \mathcal{H}_2 such that

$$(25) \quad v = \sum_n c_n (e_n \otimes f_n).$$

If the $|c_n| \neq |c_{n'}|$ for all $n \neq n'$, then the bases are unique (up to a phase $e^{i\theta}$ on each element of the basis). Note that, in general, for *arbitrary* bases $\{x_n\}$ and $\{y_m\}$ of \mathcal{H}_1 and \mathcal{H}_2 , v is expressed in general in terms of a double sum:

$$(26) \quad v = \sum_{n,m} c_{nm} (x_n \otimes y_m)$$

and compare this expression with (25).

1.2.6.3 Reduced States

⁸See Bub, Ch. 6, this volume.

1.2.6.3.a Partial Trace and the Reduced Density Operator Suppose we are given the state of a compound system, and wish to derive from it a state for one of the components. If the compound state is factorizable, then the procedure is straightforward. (The state $W = W_1 \otimes W_2$ fixes the component states to be W_1 and W_2 respectively.) But what about when it is entangled? Here we face a problem. If the state is entangled, then there is no obvious sense in which it can be ‘divided’ into a ‘part’ corresponding to one system, and a ‘part’ corresponding to the other.

The usual solution to this problem is to take the state of the component systems to be given by a *partial trace*. For any tensor-product Hilbert space, $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$, the ‘partial trace over \mathcal{H}_1 ’ is a map, $\text{tr}^{(1)}[\cdot]$, from the operators on \mathcal{H} to operators on \mathcal{H}_2 . It is the unique such map satisfying the condition that, for any density operator W on \mathcal{H} and any observable F_2 on \mathcal{H}_2 , the operator $\text{tr}^{(1)}[W]$ generates the same expectation value for F_2 as W does for $\mathbb{I}_1 \otimes F_2$ (Jauch [1968, §11-8]). The idea is that $\text{tr}^{(1)}[\cdot]$ ‘traces out’ system 1, extracting just that part of the compound state that applies to system 2. Unless W is a ‘product state’ (i.e., $W = W_1 \otimes W_2$), the reduced states derived from W are necessarily mixed states.

1.2.6.3.b Proper Versus Improper Mixtures In §1.2.3.8 I introduced the idea that a mixed state can be understood as a literal mixture of systems each in some pure state. Certainly when we are describing the state of a system chosen at random from an ensemble that was produced by literally ‘mixing’ systems in various pure states, it is quite proper to interpret the mixed state in this way. However, we now see that mixed states can arise in another way, namely, as the state of one component of a compound system that is in a non-factorizable compound state. In these cases, it is far from clear that the state (of the component) should be understood as above. Indeed, there need not even *be* an ensemble of which this component is a part. Hence mixtures that arise from taking the partial trace of the state of a compound system are normally called ‘improper mixtures’, while those that arise from a mixing of individual systems in pure states are normally called ‘proper mixtures’ (a terminology introduced by dEspagnat [1971]). Whether the probabilities generated by improper mixtures can reasonably be understood as ‘ignorance about the true pure state’ (as they can for proper mixtures) is a matter for interpretative investigation.

1.2.6.4 Correlations Compound systems that are in a non-factorizable state will exhibit correlations between the measured values of observables on the two (or more) components. Consider, for example, the statevector $v = c_1 f_1 \otimes g_1 + c_2 f_2 \otimes g_2$ (where c_1 and c_2 are non-zero coefficients), and suppose that the f_n and the g_n are eigenvectors of the observables F and G respectively. In this state, there is a correlation between the value of F on system 1 and G on system 2. Indeed, let P_{f_n} and P_{g_n} be the projections onto the subspaces spanned by f_n and g_n respectively, and let P_v be the projection onto the subspace spanned by v . Then, applying

Lüder's Rule (19), we find

$$(27) \quad \Pr_{P_v}(\mathbb{I}_1 \otimes P_{g_{n'}} | P_{f_n} \otimes \mathbb{I}_2) = \frac{\text{Tr}[(P_{f_n} \otimes \mathbb{I}_2)P_v(P_{f_n} \otimes \mathbb{I}_2)(\mathbb{I}_1 \otimes P_{g_{n'}})]}{\text{Tr}[P_v(\mathbb{I}_1 \otimes P_{f_n})]}$$

where \mathbb{I}_k is the identity on \mathcal{H}_k . Taking the trace in a basis that includes the $f_n \otimes g_{n'}$ reveals that this conditional probability is 0 when $n \neq n'$ and 1 when $n = n'$. In other words, the values of F (on system 1) and G (on system 2) are perfectly correlated.⁹ Consideration of other observables would reveal additional correlations (not always perfect correlations). We will see an example later.

1.2.7 Structure of the Space of States

We noted above (§1.2.3.8) that every convex combination of pure states is again a state. Of course, a convex combination of mixed states is (by the spectral theorem) also a convex combination of pure states, so that in fact the set of states forms a convex set (§7.1.10), a point that I shall discuss in detail later (§2.2.1). Here we note the fundamental point that the convex set of states in quantum theory is not a simplex.

This point (§7.1.10) marks a departure from classical physics, where every mixed state is uniquely decomposable in terms of pure states. One thus naturally takes the mixed state as a measure of ignorance over the pure states that appear in its decomposition. No correspondingly straightforward interpretation of mixed states in quantum theory is available, in part because the mixed states are multiply decomposable into a convex combination of pure states.

1.3 Simple Example: A Spin- $\frac{1}{2}$ Particle

An understanding of the formalism, and the issues to which it gives rise, is much aided by some experience with actual calculations, however simple. In that spirit, let us consider the example of a spin- $\frac{1}{2}$ particle. The example is well-worn, but deservedly so. While there *are* some important foundational and philosophical issues concerning quantum theory that cannot be illustrated or investigated in the context of spin- $\frac{1}{2}$ particles, many such issues can be investigated in this context.

1.3.1 Introduction of Spin into Quantum Theory

Spin was introduced in 1924 in the course of an attempt to understand the spectrum of electromagnetic radiation emitted by certain metals. In the course of that explanation, electrons were supposed to have some “two-valued quantum degree

⁹Authors will sometimes say that two observables are ‘perfectly anti-correlated’ if the two observables have the same spectrum and the value of one is always minus the value of the other. They will also occasionally reserve the term ‘perfect correlation’ for a similar case, where the value of one is always equal to the value of the other. Our use of the term ‘perfect correlation’ — according to which two observables are perfectly correlated in a state just in case the conditional probabilities for values of one, given a value of the other, are always 0 or 1 — covers both cases.

of freedom”.¹⁰ This degree of freedom was soon associated with a rotation of the electron. Because the electron is a charged body, its rotation creates a magnetic field — the electron acts as a magnet whose north and south poles lie on the axis of rotation. This magnetic property was just what was needed to explain the phenomena.

So far, the story sounds good. However, it was seen almost immediately that the rotation cannot be literal. Nonetheless, the theory of ‘spin’ was developed in the context of the new quantum theory; the name stuck, and we continue to refer to this magnetic property of electrons (and as current theory tells us, other particles) as ‘spin’.

1.3.2 Quantization of Spin

It turns out that spin is ‘quantized’, a fact already anticipated in Pauli’s characterization of the property as a ‘two-valued degree of freedom’. This fact is, classically, unexpected. To see why, consider a standard method for measuring the spin of a particle. (The method does not, in fact, work for electrons, but it illustrates the point well enough, and does work for electrically neutral particles with spin.) The relevant device is a ‘Stern-Gerlach’ device, a pair of magnets shaped and arranged to create an inhomogeneous magnetic field, that is, a magnetic field that is stronger in one direction (say, the north) than in the other. (See figure 1.)

Imagine a simple bar magnet passing between the Stern-Gerlach magnets. If the north pole points straight up so that it is close to the top magnet, then the top magnet pushes the north pole (of the bar magnet) down more than the bottom magnet pushes the south pole up, and the net result will be that the bar magnet is deflected downward. If the bar magnet enters the Stern-Gerlach magnets with the south pole facing up, then the result is the opposite: overall upward deflection. If, on the other hand, the bar magnet enters the Stern-Gerlach magnets horizontally, then it will pass straight through with no overall deflection in its path. Finally, if the bar magnet passes through neither vertically nor horizontally, then the result will be deflection, up or down, that is somewhere between the extreme cases. (The trajectories of the magnet in the two extreme cases are illustrated in figure 1.)

(One’s physical intuition might be that in the intermediate cases, the axis of rotation of the electron would snap into alignment with the magnetic field so that the intermediate cases would quickly collapse into one of the extreme cases. However, such is not the case, alas. Instead, the spinning electron would classically be expected to act as a gyroscope, thereby maintaining its original inclination with respect to the magnetic field.)

Now, imagine putting an ensemble of particles with ‘spin’ through the magnets. Rather than the classically expected result (namely, a distribution of different amounts of deflection, from ‘maximum down’ to ‘maximum up’), one finds only two results: ‘maximum down’ and ‘maximum up’ — these results are illustrated in figure 1.

¹⁰See Massimi ([2004, chs. 2,4]) for discussion.




Figure 1. An experiment involving Stern-Gerlach magnets

This result holds no matter how the Stern-Gerlach magnets are oriented. That is: notice that we could reorient the magnets so that the axis passing between them is pointing in any direction in space. Passing an electron through the device, we would, again, find that it goes either ‘up’ or ‘down’ (relative to this new direction in space). Hence we can measure the spin of a particle in any direction, and we refer to the observables thus measured as ‘spin- u ’, where u refers to some specified direction in space. These facts about spin make it clear, moreover, that classical ‘spin’ is at best a metaphor for whatever property particles with ‘spin’ have. (In any case, in non-relativistic quantum theory, the electron is normally treated as a point particle, so that its spin could not be coordinated with any spatial rotation. As it is commonly said, spin has no ‘classical analogue’.)

1.3.3 Quantum Formalism for Spin

Let us see, now, how the observables for and states of a particle with spin are represented in the formalism that I sketched above. I will consider just the degrees of freedom related to spin, ignoring, for example, spatial degrees of freedom.

1.3.3.1 Hilbert Space and Observables The Hilbert space for a single spin- $\frac{1}{2}$ particle is \mathbb{C}^2 , the space of complex column-vectors with 2 components (§7.3). The ‘observables’ of the system correspond to ‘spin’ in various directions (every

direction in space), and each will have just two possible values, which we may call ‘up’ (represented by the number $+1/2$) and ‘down’ ($-1/2$).¹¹ The spin observables in the x , y , and z directions are defined in terms of the Pauli matrices by $S_x = (1/2)\sigma_x$, and similarly for S_y and S_z . (See §7.3.1).

1.3.3.2 States The pure states can be represented by norm-1 vectors, or by projections onto the space spanned by them. Consider, for example, the statevectors

$$(28) \quad \psi = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \chi = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The vector ψ , for example, corresponds to the (pure) density operator (one-dimensional projection operator)

$$(29) \quad W = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

The vectors ψ and χ are an eigenvectors of

$$(30) \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

with eigenvalues $+1$ and -1 respectively.

Note that the expectation value of S_z in the state W is

$$(31) \quad \begin{aligned} \text{Tr} \left[\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & -\frac{1}{2} \end{pmatrix} \right] &= \text{Tr} \left[\begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{pmatrix} \right] \\ &= (1 \ 0) \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ &\quad + (0 \ 1) \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ &= \frac{1}{2} + 0 = \frac{1}{2}. \end{aligned}$$

(Recall our earlier comments about calculating traces in an appropriately chosen basis.) Of course, in general a system’s having an expectation value equal to some value, r , is not sufficient to imply that the system has the value r . (Indeed, r might not even be in the spectrum of possible values.) In this case, however, we may

¹¹ The particles that I have been discussing — those that have just two degrees of freedom (‘up’ or ‘down’) for any given direction of spin, are called ‘spin- $\frac{1}{2}$ particles’, in part because their angular momentum about any given axis is either $+\hbar/2$ (‘up’) or $-\hbar/2$ (‘down’), where \hbar is a unit of angular momentum equal, in familiar units, to 1.054×10^{-34} kg m²/s. (As is often done, I have adopted units in which $\hbar = 1$.) There are, in addition, deeper group-theoretic reasons for calling these particles ‘spin- $\frac{1}{2}$ ’ particles, but we cannot go into that point here. (There are also particles with higher spin, which means, operationally, that they have more than two degrees of freedom for each direction of spin.)

also note that the probability associated with the appropriate projection operator is 1. So, first, note that the spectral decomposition of S_z is:

$$(32) \quad S_z = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & -\frac{1}{2} \end{pmatrix} = \left(+\frac{1}{2}\right) \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \left(-\frac{1}{2}\right) \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \\ := \left(+\frac{1}{2}\right) P_{z_+} + \left(-\frac{1}{2}\right) P_{z_-} .$$

Hence the projection associated with the value $+\frac{1}{2}$ for S_z is P_{z_+} and the probability for the value $+\frac{1}{2}$ (for S_z) in the state W is

$$(33) \quad \text{Tr}[WP_{z_+}] = \text{Tr} \left[\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \right] = 1.$$

(We leave the details of the calculation to the reader. Notice that taking the trace of a matrix amounts to just adding the numbers along the diagonal. The reader might wish to prove this fact.) As I noted above in a more general context, this expression is, equivalently, the expectation value of P_{z_+} in the state W . Hence, in particular, if one agrees that ‘value r for observable F has probability 1 in state W ’ implies ‘a system in state W has value r for F ’ then we may conclude, from (33), that a system in the state W has the value $+1/2$ for S_z . (We will discuss such interpretive principles in more detail later.)

1.3.4 Incompatibility

Finally, notice that in this state, W , the expectation value of spin in the x and y directions is 0. For example,

$$(34) \quad \text{Tr}[WS_x] = \text{Tr} \left[\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix} \right] = \text{Tr} \left[\begin{pmatrix} 0 & \frac{1}{2} \\ 0 & 0 \end{pmatrix} \right] = 0.$$

This fact suggests (indeed, in this two-dimensional case, implies) that the probabilities for $S_x = +\frac{1}{2}$ and $S_x = -\frac{1}{2}$ in the state W are $\frac{1}{2}$, as we can also verify by a direct calculation. First, note that the spectral resolution of S_x is:

$$(35) \quad S_x = \begin{pmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix} = \left(\frac{1}{2}\right) \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} + \left(-\frac{1}{2}\right) \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{pmatrix} \\ := \left(+\frac{1}{2}\right) P_{x_+} + \left(-\frac{1}{2}\right) P_{x_-} .$$

As the reader may verify, $\text{Tr}[WP_{x_+}] = \text{Tr}[WP_{x_-}] = \frac{1}{2}$.

We have thus verified, in this particular case, a claim made previously made abstractly, namely, that a state that is dispersion-free (i.e., generates probabilities of just 0 or 1 for all possible values) for one observable, will necessarily not be dispersion-free for some other observables. Indeed, I said earlier that non-commuting observables that do not share eigenvectors are always incompatible, in the sense that any state that is dispersion-free on one of them is necessarily




Figure 2. Polar angles.

not dispersion-free on the other. Now notice that S_x , S_y , and S_z are mutually non-commuting, and indeed share no eigenvectors. (In this two dimensional case, non-commuting maximal observables cannot share any eigenvectors.) Hence a state that is dispersion-free for one will necessarily generate non-trivial probabilities for the others.

Indeed, consider any direction, u , in space specified relative to the z -axis by the polar angles θ and ϕ , i.e., in Cartesian coordinates, $u = (x, y, x) = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$. (See Figure 2.) Then the associated spin observable is represented by the matrix

$$(36) \quad S_u = \frac{1}{2} \begin{pmatrix} \cos \theta & e^{-i\phi} \sin \theta \\ e^{i\phi} \sin \theta & -\cos \theta \end{pmatrix}.$$

(One reasonable and quick justification of this expression is to note that $S_u = S_x \sin \theta \cos \phi + S_y \sin \theta \sin \phi + S_z \cos \theta$.) The only pairs of such operators that commute are anti-parallel; i.e., they correspond to spin in anti-parallel directions (and such operators are just multiples of one another by a factor of -1).

(One should keep in mind, however, that Gleason’s theorem does not hold for our 2-dimensional space. Hence the density operators do not define all states, in this case. Indeed, Bell [1964] shows how to define a dispersion-free measure over the projections on \mathbb{C}^2 in terms of an additional ‘hidden’ parameter. Moreover, the quantum-mechanical states are obtainable by averaging over the possible values of the hidden parameters with an appropriate probability distribution over them.)

1.3.5 The Bloch Sphere

The Hilbert space \mathbb{C}^2 is used to represent any two-level quantum system, and such systems are of great interest in quantum theory, all the more so in recent years, as increasing interest in quantum information and quantum computation has focused

attention even more on such systems (because they are the quantum analog of a classical ‘bit’ — see Bub, Ch. 6, this volume). A careful study of the pure states on \mathbb{C}^2 is often aided by the representation of those states in terms of the Bloch sphere. Note that any pure state on \mathbb{C}^2 can be represented by a vector of the form $v = \cos(\theta/2)\psi + e^{i\phi} \sin(\theta/2)\chi$ (using the notation of equations 28).¹² Hence, again referring to figure 2, we can represent each distinct pure state as a unique point on the surface of a unit sphere (in \mathbb{R}^3), normally called the ‘Bloch sphere’. The ‘north pole’ of the sphere corresponds to the state ψ and the ‘south pole’ to the state χ .

In fact, however, the ‘Bloch sphere’ is a ball. The interior points correspond to mixed states, as follows. Every density operator, W , on \mathbb{C}^2 can be written as

$$(37) \quad W = \frac{\mathbb{I} + \vec{r} \cdot \vec{\sigma}}{2}$$

for $\vec{\sigma}$ the ‘vector’ of Pauli matrices (§7.3.1) and \vec{r} a vector from \mathbb{R}^3 with $\|\vec{r}\| \leq 1$. The components of \vec{r} determine a point inside the Bloch sphere representing the corresponding density operator. (Note, in particular, that $\vec{r} = (0, 0, 1)$ corresponds to the pure state given by $\theta = 0$, as it should.)

1.4 Dirac Notation

I will return to the example of a spin- $\frac{1}{2}$ particle later to illustrate a number of issues in quantum theory. When I do so — and, indeed, throughout the remainder of this essay — it will be helpful to have at hand a useful notation, the so-called ‘Dirac bra-ket’ notation, used commonly by both physicists and philosophers.

1.4.1 Bras and Kets

In the bra-ket notation, vectors are denoted by (and sometimes called) ‘kets’, $|v\rangle$. In the discussion above, for example, the column vector ψ in (28) might be denoted $|z_+\rangle$. Elements of the dual space (the ‘row vectors’ in our discussion above — see §7.1.8) are denoted by ‘bras’, $\langle v|$. In our example above, there is a natural 1-1 map from the kets (column vectors) to the bras (row vectors):

$$(38) \quad \begin{pmatrix} a \\ b \end{pmatrix} \rightarrow (a^* \quad b^*).$$

The bras thus define (continuous) linear functionals in the obvious way. Letting

$$(39) \quad |v\rangle = \begin{pmatrix} a \\ b \end{pmatrix} \quad \text{and} \quad |w\rangle = \begin{pmatrix} c \\ d \end{pmatrix},$$

¹²The claim is not that every vector can be written in this form, but that every pure state can be represented in this form. Recall that an overall phase factor does not affect the probabilities generated by a vector. Hence we may assume, without loss of generality, that the coefficient of ψ is real.

the linear functional (bra) $\langle v|$ acting on the vector (ket) $|w\rangle$ is

$$(40) \quad (a^* \quad b^*) \begin{pmatrix} c \\ d \end{pmatrix} = a^*c + b^*d$$

and is written, in the Dirac notation, as (the ‘bra-ket’) $\langle v|w\rangle$. (The reader might wish to check that the functional thus defined is indeed linear.) Of course, as it must be, $\langle v|w\rangle$ is also the inner product of $|v\rangle$ with $|w\rangle$, given (38). (In this notation, we continue to write $\|v\|$ for the norm of a vector, instead of $\| |v\rangle \|$.)

In the general case, i.e., where \mathcal{H} is any (complex) Hilbert space (countable-dimensional at most), we take the elements of \mathcal{H} to be kets, and the elements of the dual space \mathcal{H}^* to be bras. Inner products may now be written $\langle v|w\rangle$, which denotes *both* the linear functional $|v\rangle$ acting on the vector $|w\rangle$ *and* the inner product of the vectors $|v\rangle$ and $|w\rangle$.

1.4.2 Operators

The operator, F , acting on the vector $|v\rangle$ is written $F|v\rangle$. The expectation value of the observable F in the state $|v\rangle$ is written $\langle v|F|v\rangle$, which is notationally (and numerically) equivalent to $\langle v|Fv\rangle$, the latter to be read as the inner product of $|v\rangle$ with the vector $F|v\rangle$. The expression $\langle w|F|v\rangle$ is defined similarly.

Corresponding to what is sometimes called the ‘vector direct product’

$$(41) \quad \begin{pmatrix} a \\ b \end{pmatrix} (c \quad d) = \begin{pmatrix} ac & ad \\ bc & bd \end{pmatrix},$$

we can define $|v\rangle\langle w|$ to be the operator on \mathcal{H} defined by

$$(42) \quad (|v\rangle\langle w|)|x\rangle = \langle w|x\rangle |v\rangle.$$

Notice that simple symbol-manipulation would generate the same result.

1.4.3 Using the Dirac Notation

As I just hinted, the Dirac notation is enormously useful, once its true meaning is understood, and dangerous otherwise. Its power — and danger — lies in the fact that it allows one more or less to ignore various distinctions, such as the distinction between a vector and a linear functional (element of a dual space). It also can be very helpful for ‘coordinate-free’ calculations. For example, we can discuss the theory of spin- $\frac{1}{2}$ particles without bothering with Pauli matrices and so on. Consider the basis $\{|z_+\rangle, |z_-\rangle\}$ for \mathbb{C}^2 , where $|z_+\rangle$ is the state that assigns probability 1 to the value $+\frac{1}{2}$ for S_z and so on — note that we do not need to worry about how to represent this state as a column of complex numbers. It is sufficient to carry out calculations to note that for a direction in space, u , specified by the angles θ and ϕ relative to the z -axis:

$$(43) \quad |u_+\rangle = \cos\left(\frac{\theta}{2}\right) e^{-i\frac{\phi}{2}} |z_+\rangle + \sin\left(\frac{\theta}{2}\right) e^{i\frac{\phi}{2}} |z_-\rangle$$

$$(44) \quad |u_-\rangle = -\sin\left(\frac{\theta}{2}\right) e^{-i\frac{\phi}{2}} |z_+\rangle + \cos\left(\frac{\theta}{2}\right) e^{i\frac{\phi}{2}} |z_-\rangle.$$

The spin observables are then represented by

$$(45) \quad S_u = \frac{1}{2}|u_+\rangle\langle u_+| - \frac{1}{2}|u_-\rangle\langle u_-|.$$

Note, for example, that $\langle z_+|u_+\rangle = \cos\left(\frac{\theta}{2}\right)e^{-i\frac{\phi}{2}}$ and $\langle z_-|u_+\rangle = \sin\left(\frac{\theta}{2}\right)e^{i\frac{\phi}{2}}$, facts that are immediately read off of (43). Hence, for example, the probability that a system in the state $W = |z_+\rangle\langle z_+|$ has the value $+\frac{1}{2}$ for the observable S_u can be quickly calculated as

$$(46) \quad \text{Tr}[|z_+\rangle\langle z_+|(|u_+\rangle\langle u_+|)]$$

$$(47) \quad = \langle z_+|(|u_+\rangle\langle u_+|)|z_+\rangle$$

$$(48) \quad = \langle z_+|u_+\rangle\langle u_+|z_+\rangle$$

$$(49) \quad = |\langle z_+|u_+\rangle|^2$$

$$(50) \quad = \cos\left(\frac{\theta}{2}\right)^2.$$

(To get from the first to the second line, calculate the trace using the basis $\{|z_+\rangle, |z_-\rangle\}$.) The genius of Dirac's notation is that one can, as illustrated here, simply 'do the symbolically natural thing' and get the correct answer. For example, the third line follows from the second by 'erasing the parentheses and joining the bars'. Conceptually, we allowed the operator $|u_+\rangle\langle u_+|$ to act on $|z_+\rangle$, obtaining the vector $\langle u_+|z_+\rangle|u_+\rangle$, then took the inner product of this vector with $|z_+\rangle$ (or, applied the linear functional $\langle z_+|$ to $\langle u_+|z_+\rangle|u_+\rangle$). The convenience of the notation can also, however, lead one to forget conceptually important distinctions.

Keep in mind, moreover, that the convenience of not having to worry about explicit (e.g., matrix) representations of vectors and observables can also lead one to write down some rather silly, or at least physically opaque, states. One frequently, for example, sees 'states' such as $|\text{cat dead}\rangle$ or $|\text{Sarah sees the pointer}\rangle$ written down. The Dirac notation naturally tempts one to write down such expressions, but we are so far from knowing whether such 'states' correspond to some pure vector state, and if so, what their properties are, that such expressions are best left to cartoons.

1.5 Transformations

We have now seen how to represent observables, and how to calculate expectation values (and probabilities). While such matters are indeed at the heart of the theory, there are other aspects of the formalism that are important for philosophical and foundational discussions. In particular, this subsection discusses transformations, both of the states of physical systems and of the observables associated with those systems. Along the way, I will have occasion to mention some theorems that are fundamental for the foundations of quantum mechanics.

1.5.1 Groups and Their Representations

1.5.1.1 Motivation Galileo observed that the laws of motion do not depend on the constant velocity of the ‘lab’ (frame of reference) in which they are applied. (For example, in the hull of a ship moving with constant velocity — more precisely, moving inertially — “jumping with your feet together, you pass equal spaces in every direction”, as Galileo writes, just as you would back on shore.) Neither do they depend on one’s location, nor on the time at which they are applied, nor on the direction in which one is facing. In other words, the laws are invariant under certain transformations, namely, boosts (changes in velocity), spatial translations, temporal translations, and rotations. These sorts of transformation are represented, mathematically, by groups, and in the case of the ‘Galilean transformations’ that I just mentioned, the group is normally called the ‘Galilean group’.¹³ Hence group theory (§7.6) is the natural context in which to study, among other things, the ‘invariances’ of quantum theory.

The motivation here is that the properties of a group are exactly the properties normally thought to apply to ‘invariance transformations’. In particular, if α and β are transformations that each individually leave the laws unchanged, then the composition of α followed by β is also such a transformation. Similarly, if α is a such transformation, then there is the transformation that ‘undoes’ what α did, that is the inverse of α . Notice, for example, that the composition of two Galilean transformations is another one, and that each transformation has an inverse.¹⁴

Groups show up in other contexts as well. Suppose, for example, that we are interested (as we soon will be) in the dynamics of a closed physical system. One way to think about the time-evolution of the state of a system is as a transformation on the set of states. The set of all such time-evolutions, then, plausibly should form a group. The identity represents ‘no change’ (or the degenerate case of evolution over no time). The product represents one period of evolution followed by another. And the inverse represents ‘reversed’ evolution, or evolution backwards in time. (If a given theory is not time-reversible, then we would be dealing with a semi-group rather than a group.)

Now, often one specifies a group abstractly, that is, by specifying the products and inverses in the group without representing it as a group of transformations on some set (such as the set of physical states of a system). The most trivial example is the group Z_2 , which contains two elements, x and y . The multiplication rule is: $xy = x$, $yx = x$, $xx = y$, and $yy = y$. The identity is (clearly) y , while x and y are their own inverses. Notice that we specified this group without referring to any

¹³More precisely, the Galilean group is $(\mathcal{R} \ltimes \mathcal{V}) \ltimes (\mathcal{A} \times \mathcal{T})$, where \times is the direct product, \ltimes is the semi-direct product, and \mathcal{T} , \mathcal{A} , \mathcal{V} , and \mathcal{R} are the (sub-)groups of temporal translations, spatial translations, boosts, and rotations, respectively (§7.6.2). If the Galilean group is defined, first and foremost, as the set of affine (parallel-line-preserving) maps from E , the Euclidean 4-dimensional manifold (§7.5.2) of events (space-time), to itself that preserve simultaneity of events and the distance between simultaneous events, then it turns out that the subgroups mentioned above are not all normal, as implied by the use of semi-direct products where one might expect direct products.

¹⁴See Brading and Castellani, Ch. 13, this volume, for more nuanced discussion.

specific mathematical objects — the symbols ‘ x ’ and ‘ y ’ are just names for the two elements of this group and by themselves have no further mathematical content. But we could also ‘represent’ the group Z_2 as, for example, the group of maps from any two-element set to itself, with y being the identity map, and x being the map that swaps the elements (maps each to the other). (Another representation of Z_2 takes x to be complex conjugation, $*$, and y to be $**$.)

1.5.1.2 Wigner’s Theorem Thinking of groups as ‘collections of symmetry transformations’, the very idea that these transformations are ‘symmetries’ suggests that they should not change the relationships amongst states. In particular, a symmetry transformation on the space of states should be such that a system in state $|\psi\rangle$ generates the same probabilities for observables both before and after the transformation (at least for observables that are supposed to be invariant under this symmetry, or have been ‘transformed along’ with $|\psi\rangle$, in the sense that their eigenvectors are also transformed). How might such transformations be represented?

Notice that a unitary operator (§7.2.6) fits the bill very nicely. Indeed, we *define* a unitary operator as, in part, one that preserves inner products. There is an important near-converse to this fact, due to Wigner ([1931, p. 251]).

Theorem (Wigner): Let \mathcal{H} be a Hilbert space over \mathbb{C} and let $T : \mathcal{H} \rightarrow \mathcal{H}$ be a 1-1 (but not necessarily linear) map satisfying $\langle Tw|Tv \rangle = \langle w|v \rangle$ for any $|w\rangle, |v\rangle \in \mathcal{H}$. Then

$$(51) \quad T|v\rangle = \varphi(v)U|v\rangle$$

where U is either unitary or anti-unitary and $\varphi(\cdot)$ is a ‘phase function’, a complex-valued function on \mathcal{H} whose values have modulus 1.

(Any anti-unitary operator, T , can be written as $T = UK$, where K is the ‘complex conjugation’ operator. Hence the anti-unitary transformations are just the unitary ones, followed by complex conjugation. Time-reversal, for example, is often associated with complex conjugation.)

One normally rules out the anti-unitary case on various grounds related to the ‘unphysical’ nature of such transformations; in particular, they are not continuously connected to the identity. In order to make this notion precise, one would need to introduce a topology on the group. In the typical cases of interest, the group is continuously parametrized (§7.6.4) by some set of real indices so that the group in fact forms a manifold (§7.5.2); i.e., it is a Lie group (§7.6.5). In these cases, a topology is already given. The significance of being continuously connected to the identity is just that in this case, one has the picture of the group transformations being built up from transformations that are ‘infinitesimal’, i.e., ‘as close as you like to doing nothing at all to the system’ (the identity transformation). Of course, if we are talking just about symmetries, there is no reason to suppose that being continuously connected to the identity is a necessary condition — consider, just to mention the most obvious examples, time-reversal, or spatial reflection. On the other hand, if the symmetries in question are supposed to correspond, ultimately, to actual physical processes (such as dynamical evolution of a

closed system), then continuous connectedness to the identity begins to look more compelling.

Hence, in general, symmetries in quantum theory are represented in terms of these maps, T , with U unitary or anti-unitary, and often under the assumption (or hope) that U is unitary.

1.5.1.3 Projective Representations In the expression (51) one not only (normally) sets aside the case where U is anti-unitary, but also (normally) seeks maps, T , such that $\varphi(v)$ is identically 1. In this case, the representation of the symmetry group is just given in terms of a group of unitary operators. Such representations are particularly nice because much is known about unitary operators. (See §1.5.1.4 for an important example.) But one is not always so fortunate as to be able to find this sort of representation, often called a ‘unitary’ or ‘ordinary’ representation (§7.6.8). Sometimes one must live with the phase function’s being non-trivial. In this case, the representation is called ‘projective’.

The reason is as follows. Let \mathcal{H} be a Hilbert space, and consider the set, $P\mathcal{H}$, of equivalence classes of vectors from \mathcal{H} , where two vectors are equivalent if and only if they lie in the same one-dimensional subspace. $P\mathcal{H}$ is a *projective Hilbert space*, whose structure is given by the ‘angles’ between the rays of \mathcal{H} (the modulus of the inner product of normalized representatives from the rays). When the phase function in (51) is non-trivial, the resulting transformation still generates an automorphism of $P\mathcal{H}$. (Moreover, we have already observed that the pure states in quantum theory can, for the purposes of calculating probabilities, be just as well represented by one-dimensional projections as by statevectors. Hence it should come as no surprise that projective representations of a group can still preserve all probabilities.) Hence, while ordinary representations tend to be easier to handle, there is nothing terribly inconvenient or problematic about projective representations, and one is sometimes forced to use them.

1.5.1.4 Stone’s Theorem Unitary representations are particularly nice, because they can be ‘generated’ by self-adjoint operators. Note, first, that given any self-adjoint operator, F , the operator e^{iF} is unitary. Moreover, the family of operators $e^{i\alpha F}$ with α a real parameter forms a continuously parametrized group of unitary operators, where $e^{i\alpha F} e^{i\alpha' F} = e^{i(\alpha+\alpha')F}$. (Note that $\lim_{\alpha \rightarrow 0} e^{i\alpha F} = I$, i.e., this group is continuously connected to the identity.) Now suppose that we are interested in representing a continuously parametrized group, \mathcal{G} , as a family of unitary operators on a Hilbert space. Because of the nice behavior of the $e^{i\alpha F}$, one would very much like to find an F that generates a representation of \mathcal{G} . We are in luck:

Theorem (Stone [1932]): Let U_α be a (weakly) continuous unitary representation of \mathcal{G} on a Hilbert space, \mathcal{H} .¹⁵ Then there exists a self-adjoint operator, F , on \mathcal{H} such that $U_\alpha = e^{-iF\alpha}$.

¹⁵ The map $\alpha \mapsto U_\alpha$ is *weakly continuous* if and only if $\langle w|U_\alpha|v \rangle$ is a continuous function of α for every $|v\rangle, |w\rangle \in \mathcal{H}$. Indeed, Stone’s theorem holds under weaker conditions when \mathcal{H}

Because so many of the groups of interest in non-relativistic (and indeed relativistic) quantum mechanics have the requisite properties, Stone's theorem is of fundamental importance for the theory. We shall see some examples of its use later.

1.5.2 Dynamics

1.5.2.1 Some Initial Assumptions about Dynamical Evolution Dynamical evolution of the state of a system is just a kind of transformation on the space of states. I begin with a few simplifying assumptions, which will turn out to be sufficient to determine the form of dynamical evolution.

Keeping in mind that density operators are mixed states, and hence linear combinations of pure states, let us assume that the evolution of a density operator is induced by the evolution of the pure states of which it is composed. The idea here is that a density operator can represent simple physical mixing, and in that case at least, it should evolve as described. Suppose, for example, that we have a mixture, in proportions r and $1 - r$ (with $0 < r < 1$), of two different types of system, the first in the pure state P_1 and the second in the pure state P_2 . The corresponding density operator is $rP_1 + (1 - r)P_2$. If the systems evolve without interacting with one another (for example, they might be physically isolated from one another), then one would expect that if the systems in P_n evolve to the state P'_n , the mixture evolves to $rP'_1 + (1 - r)P'_2$, or so I shall assume. In that case, we can focus our attention on the pure states, and hence on the (normalized) vectors of a Hilbert space.

Notice that this argument certainly does not apply to density operators that arise from partial tracing (i.e., 'reduced density operators'). Indeed, in general, such operators will not evolve in the way described. But rather than determining their evolution directly, one may derive it from evolution of the state of the compound system of which it is a component — and if that compound system is in an improper mixture, then repeat the procedure.

We are therefore primarily interested in the dynamics of an isolated physical system (though certainly the system would in general experience internal interactions), represented by a pure state. The question then becomes: which transformations on a Hilbert space for this system are possible dynamical evolutions of the state of the system?

Symmetry is again a helpful tool. Let $T : \mathcal{H} \rightarrow \mathcal{H}$ be a map from the Hilbert space \mathcal{H} to itself representing the time-evolution (over some given stretch of time) of a closed system. Because the system is closed, it seems reasonable to suppose that this T should be a symmetry, in the sense that we have already noticed: $|\langle v|w\rangle|^2 = |\langle Tv|Tw\rangle|^2$ for all $|v\rangle, |w\rangle \in \mathcal{H}$. (Of course, the ultimate justification for this supposition is empirical success.) Hence (51) applies. We will assume, further, that the phase function $\varphi(v)$ is identically 1, i.e., that time-evolution is given by an

has a countable basis (i.e., it is 'separable'), in which case the functions $\langle w|U_\alpha|v\rangle$ need only be Lebesgue-measurable. See Riesz and Sz.-Nagy ([1955, §137]).

ordinary representation of some group of evolution operators. Finally, recall that the anti-unitary operators are not continuously connected to the identity, meaning that, in this case, if they were used to represent time-evolution, there would be no way to represent evolution over infinitesimal times. Under the assumption that time is continuous, we are led to suppose that time-evolution is given by some group of unitary operators.

Notice, now, that our original argument for the conclusion that the evolution of mixed states should follow from the evolution of pure states, essentially by the assumption of linearity, is encapsulated in this result. As a transformation on \mathcal{H} , a unitary operator U also generates a transformation on the operators on \mathcal{H} , and assuming that dynamical evolution ought not change the relationship (that is, the relations definable purely in terms of the structure of the Hilbert space) between pure states and the mixtures that they compose, we must use U to generate the transformation on mixed states by $W \mapsto U^{-1}WU$. Indeed, this expression is the most general form of the standard dynamics of a closed system in quantum mechanics.

To see where it comes from, consider the density operator, $W(t) = \sum_n w_n(t)P_n(t)$. Letting $\{|\psi_{n,i}(t)\rangle\}$ be an orthonormal basis determined by $W(t)$ (where the index n ranges over the spectral projections $P_n(t)$ and the index i ranges over the dimension of $P_n(t)$), we may write

$$(52) \quad W = \sum_{n,i} w_{n,i}(t) |\psi_{n,i}(t)\rangle \langle \psi_{n,i}(t)|.$$

According to our earlier assumption, the evolution of W will be given in terms of the evolution of the $|\psi_{n,i}(t)\rangle$, which means in particular that the coefficients $w_{n,i}(t)$ will be time-independent. (Keep in mind, here, that we presume W to be the state of a closed (isolated) system. Without that assumption, nothing said here would be at all plausible.) Hence, if U_t is the evolution operator for the system, we may simply apply it to the summands:

$$(53) \quad W(t) = \sum_{n,i} w_{n,i} \left(U_t |\psi_{n,i}(0)\rangle \right) \left(\langle \psi_{n,i}(0) | U_t^* \right)$$

$$(54) \quad = U_t \left(\sum_{n,i} w_{n,i} |\psi_{n,i}(0)\rangle \langle \psi_{n,i}(0)| \right) U_t^*$$

$$(55) \quad = U_t W(0) U_t^*,$$

where we have used the linearity of the U_t . Finally, recall that for any unitary operator, $U, U^* = U^{-1}$.

Finally, it is worth noticing that for any given unitary map, U , on \mathcal{H} , considered to be a ‘symmetry’, the map $F \mapsto U^{-1}FU$ for all operators F is the ‘correct’ corresponding symmetry of operators, at least in the sense that for any $|v\rangle \in \mathcal{H}$ and any operator, F , on \mathcal{H} , $U(F|v\rangle) = (UFU^{-1})U|v\rangle$. That is, one can either ‘apply the operator F to the vector $|v\rangle$, then transform according to U ’, or ‘transform

according to U then apply the transformed operator to the transformed vector', and in both cases the result is the same.

1.5.2.2 The Hamiltonian We can say, then, that the evolutions on \mathcal{H} are generated by unitary operators. But which ones? For example, which operator represents the evolution of a free particle? Which operator represents the evolution of a particle under the influence of some given potential energy? Some progress on that question is made via Stone's Theorem.

Above I claimed that dynamical evolution has all the properties of a group. In particular, let $U_{1,2}$ represent the evolution of a system from time t_1 to time t_2 , and similarly for $U_{2,3}$. Then it would seem that evolving from t_1 to t_2 , and thence to t_3 , is the same as evolving from t_1 to t_3 ; in other words, $U_{1,3} = U_{2,3}U_{1,2}$. (Keep in mind that we are thinking of these $U_{m,n}$ as operating on some space of states — hence the ordering.)

A slightly stronger, but still quite compelling, assumption is that of 'time homogeneity'. Imagine a system evolving under the influence of some time-independent constraints (for example, time-independent potential energies). Then, if $t_3 - t_2 = t_2 - t_1$, the evolution operator $U_{1,2}$ should in fact be the same as $U_{2,3}$. (Keep in mind that these operators transform the entire space; we are not assuming that a given single system will 'do the same thing' from t_2 to t_3 as it did from t_1 to t_2 , but rather that two different isolated systems in the same state at two different times will do the same thing for the next equal stretch of time.) In this case, evolution operators require just a single parameter indicating the length of the time interval concerned, and we then have the relation: $U_t U_{t'} = U_{t+t'}$. We assume, as well, that U_0 is the identity, meaning that 'nothing happens instantaneously'.

Notice the similarity to an additive group. Indeed, a set of operators, U_t , obeying this rule forms a semi-group. But there is, as well, reasonable motivation for the existence of inverses in this case, namely, that they correspond to reverse evolution (evolution backwards in time). Writing the time-reversal of U_t as U_{-t} then we require that $U_t U_{-t} = U_{-t} U_t = \mathbb{I}$. In that case, the U_t form a one-parameter group.

Finally, we will add an assumption of continuity in time. In particular, we will assume the 'weak continuity' of the group U_t (see footnote 15). We can then apply Stone's Theorem to learn that, for *any* group representing time-evolution of a quantum system, there is some self-adjoint operator, H , such that $U_t = e^{-iHt}$.

How do we know which operator H to choose for a given system? Alas, the usual answer to this question is via 'quantization' of the Hamiltonian of an analogous classical system. For such a system, the Hamiltonian is usually the classical total energy. I say 'Alas', because although at a practical level quantization is *normally* straightforward, there is, as yet, no completely satisfactory foundational account of the nature of the 'analogy' between a given quantum system and its classical 'analogue'.¹⁶ The most frequent examples are, however, straightforward. For

¹⁶See Landsman, Ch. 5, §4, this volume.

example, the classical kinetic energy of a particle moving in one dimension is $p^2/2m$ (where p is the classical momentum and m is the mass of the particle), and the quantum-theoretic ('quantized') Hamiltonian is $P^2/2m$ (where P is the momentum operator). Hence the time evolution of a free particle in quantum theory is given by $|\psi(t)\rangle = e^{-iP^2t/2m}|\psi(0)\rangle$.

1.5.2.3 Equations of Motion In this final section about the formalism of quantum theory, I briefly introduce the standard equations of motion in non-relativistic quantum mechanics. Much of the actual practice of quantum mechanics consists in solving these equations, either exactly (in rare cases where analytic solutions are obtainable), or approximately (most of the time, either using standard techniques from perturbation theory, or using numerical approximations).

1.5.2.3.a The Schrödinger Equation Consider the group of time-evolutions $U_t = e^{-iHt}$. These U_t uniquely solve the differential equation

$$(56) \quad \frac{\partial U_t}{\partial t} = -iHU_t.$$

But $\frac{\partial U_t}{\partial t}|\psi(0)\rangle$ is just the time-derivative of $|\psi(t)\rangle$ at time t , so that, allowing the operators on both sides of (56) to act on $|\psi(0)\rangle$:

$$(57) \quad \frac{\partial}{\partial t}|\psi(t)\rangle = -iH|\psi(t)\rangle.$$

Equation (57) is the *Schrödinger equation* for a time-independent Hamiltonian. A similar equation holds for 'bra' vectors:

$$(58) \quad \frac{\partial}{\partial t}\langle\psi(t)| = \langle\psi(t)|(iH).$$

(Recall our earlier comments about the simultaneous usefulness and seductiveness of the Dirac notation. The reader is invited, here, to consider what the terms in this equation *really* represent.) Of course, the evolution of mixed states is still given by (53).

If H depends on time, then we can still consider each infinitesimal evolution (from t to $t + dt$) to be given by unitary operators $e^{-iH(t)}$. In general, it is non-trivial to build up finite-time evolution operators from these infinitesimal ones. But when the Hamiltonian is time-independent, then of course we may define $U_t = e^{-iHt}$. If, moreover, the system is in a 'stationary state' at time $t = 0$ — that is if it is in an energy eigenstate, $|\psi_E(0)\rangle$, of fixed energy, E , i.e., $H|\psi_E(0)\rangle = E|\psi_E(0)\rangle$ — then the evolution takes the simple form $|\psi_E(t)\rangle = e^{-iHt}|\psi_E(0)\rangle = e^{-iEt}|\psi_E(0)\rangle$. That is, the system remains in the same one-dimensional subspace; only the phase, e^{-iEt} , changes with time.

Of course, (57) tells us how expectation values change as well. A straightforward application shows that for any operator, F , (57) implies:

$$(59) \quad \frac{d}{dt}\langle F \rangle = -i\langle [F, H] \rangle + \left\langle \frac{\partial F}{\partial t} \right\rangle,$$

where $\langle \cdot \rangle$ is the expectation value of the operator on the inside in some (here unspecified) state and H is the Hamiltonian.

Such is the standard account of ‘Schrödinger evolution’. One should note, however, that in practice the system is often *not* isolated, so that it is necessary to add potentials to the Hamiltonian that represent the influence of some external system, σ_{ext} , on the system of interest, σ . The problem here is that it is often impractical to attempt to model the entire compound system (σ_{ext} and σ) whereas one has at least a fighting chance to model the influence of σ_{ext} on σ as an external potential applied to σ . A typical example involves nano-electronics, where, for example, one might be interested in studying a potential difference applied across a molecule. In principle, one would include the electrodes at either end of the molecule. In practice, just modeling the molecule is already very difficult, and including the electrodes in the system is completely unfeasible. Instead, one simply encodes their effect on the molecule into the potential term in the Hamiltonian. In general, doing so results in *non-unitary* evolution (because the system is not closed).

1.5.2.3.b The Schrödinger and Heisenberg Pictures We have been thinking about dynamics in terms of the evolution of states. One may equivalently consider that the state is constant in time, but evolve the observables instead. Indeed, suppose that the state of the system is $|\psi(t)\rangle$, evolving under the unitary evolution U_t . Then the expectation value of the observable F at time t is $\langle \psi(0)|U_t^* F(U_t|\psi(0))\rangle$. So we may just as well let the state be constant in time, i.e. $|\psi(t)\rangle = |\psi(0)\rangle$, while supposing that the *observables* change according to $F(t) = U_t^* F(0)U_t$. Clearly, the expression for the expectation value will be the same in either case. Hence the two pictures are empirically equivalent.

The first picture (where states evolve in time and observables are constant) is normally called the ‘Schrödinger picture’, while the second picture (where states are constant and observables evolve in time) is normally called the ‘Heisenberg picture’.¹⁷

1.5.2.3.c The Heisenberg Equation In the Heisenberg picture, how do the observables change in time? We will (temporarily) index ‘Schrödinger’ observables with an ‘S’ and the corresponding ‘Heisenberg’ observables with an ‘H’. We will assume, as well, that the Schrödinger observables do not depend explicitly on time (as is normally the case in basic applications). At any finite time, then, we would have $F^{\text{H}}(t) = U_t^* F^{\text{S}}U_t$. (See §1.5.2.3.b.) Then

$$(60) \quad \frac{dF^{\text{H}}}{dt} = \frac{\partial U_t^*}{\partial t} F^{\text{S}}U_t + U_t^* F^{\text{S}} \frac{\partial U}{\partial t}$$

¹⁷There is a third picture, the ‘interaction picture’ (sometimes called the ‘Dirac picture’), that combines the first two. In this picture, both the states and the observables evolve in time. The evolution of a system due to the free part of its Hamiltonian is encoded into the evolution of the state, and the evolution of a system due to ‘the rest’ of the Hamiltonian (the ‘interaction’ part of the Hamiltonian) is encoded in the evolution of the observables.

$$(61) \quad = iU_t^* H^S U_t U_t^* F^S U_t - iU_t^* F^S U_t U_t^* H^S U_t$$

$$(62) \quad = i[U_t^* H^S U_t, F^H]$$

$$(63) \quad = i[H, F^H]$$

where we have first used (56) and inserted $U_t U_t^*$ into each term, and then we used the fact that H^S commutes with U_t , hence $U_t^* H^S U_t = H^S = H^H$. (We have therefore also dropped the index from the Hamiltonian in (63).) This equation is normally called the ‘Heisenberg equation’.¹⁸

There is an important similarity between the Heisenberg equation of motion and the equation of motion for a classical observable (function on phase space) $f(x, p)$, written in terms of the Poisson bracket,

$$(64) \quad \{f, g\} = \frac{\partial f}{\partial x} \frac{\partial g}{\partial p} - \frac{\partial f}{\partial p} \frac{\partial g}{\partial x}.$$

The commutator and the Poisson bracket have similar algebraic properties, and moreover, the classical equation of motion is given by $\frac{d}{dt} f = \{H, f\}$, where H is the classical Hamiltonian. Hence one way of thinking about the relationship between classical and quantum theory is in terms of an algebraic analogy between commutators and Poisson brackets.¹⁹

1.6 Preview of Philosophical Issues

Philosophy of non-relativistic quantum theory has traditionally been largely concerned with four issues (which is not to suggest, of course, that there are not other issues that have been discussed).

The first issue (discussed in §2) is in some ways the most fundamental, but also, in some circles, the least discussed, namely, the origins or justification of the quantum-theoretic formalism. How do we find ourselves using Hilbert spaces and self-adjoint operators on them rather than phase spaces and functions on them? Given that the other issues facing the philosopher of quantum theory are mirrored in certain features of the formalism, answering (or at least addressing) this question could prove very useful. There are some reasonably illuminating attempts to explain why quantum theory uses the formalism it does. The next section is devoted to a discussion of a few such attempts.

The second issue (discussed in §3) concerns empirical content. A close reading of, for example, some of the founders of quantum theory (especially Niels Bohr)

¹⁸There is a similar equation for the evolution of the density operator in the Schrödinger picture: $\frac{d}{dt} W = -i[H, W]$. Its derivation is similar to the one given above. Notice the change of sign. There is no contradiction here. This equation describes the time-evolution of density operators (states) in the Schrödinger picture. Equation (63) describes the time-evolution of operators (observables) in the Heisenberg picture.

¹⁹Dirac proposed to *define* quantization in terms of this analogy. It turns out that the analogy cannot be carried through rigorously, at least not in its original form. For further discussion of quantization and this algebraic analogy, see Landsman, Ch. 5, this volume, especially §4. For further discussion of Poisson manifolds in classical mechanics (hence, the classical side of the analogy), see Butterfield, Ch. 1, this volume, especially §5.

reveals a deep concern on their part with the question how the quantum formalism gets its empirical meaning. I will sketch a proposal, motivated by Bohr, for how to understand the connection between formalism and observation.

The issue of empirical content is closely connected with the issue of incompatibility, and more precisely, the uncertainty relations (discussed in §4). Why is it not possible to determine precise values for all observables? Does it follow from a deeper claim that one observable's having a definite value can preclude another, incompatible, observable from doing so? Some of the founders of quantum theory (again, especially Bohr) were concerned to explain how it is that the conditions required for the empirical well-definedness of some pieces of the formalism somehow preclude the simultaneous satisfaction of the conditions required for the empirical well-definedness of other pieces of the formalism. Others (especially Einstein) were concerned, instead, to 'beat' the uncertainty principle, initially by arguing that one *can* determine precise values for incompatible observables, and later by arguing for the weaker claim (in the famous Einstein-Podolsky-Rosen paper) that they must *have* precise values (simultaneously), even if we cannot determine (for example, measure) them.

The fourth issue (discussed in §5), again related to the previous one, arises from Gleason's Theorem: if there are no dispersion-free (sometimes in this context called 'two-valued') states in quantum theory, then apparently some observables can sometimes be in 'indeterminate' states. Even worse, it is easy to describe physically plausible scenarios in which observables that one normally believes to be definite will *not* be assigned a definite value by the state assigned via quantum theory, under a standard interpretation of the relationship between states and values for observables. Famously, Schrödinger described such a scenario, in which a cat was somehow 'neither alive nor dead'. The obvious way to avoid this problem — by interpreting the probabilities generated by the quantum state epistemically (i.e., as expressions of ignorance about the actual values, as classical probabilities often are) seems to run into serious philosophical, indeed logical, difficulties. Hence other solutions are proposed, in the form of sometimes quite exotic interpretations of the formalism. While many of these solutions can be seen to avoid the basic problem, each has its serious pitfalls as well — and therefore, detractors. No generally accepted response to the problem seems to be on offer.

The final major issue (discussed in §6) is the non-locality of the theory, as implied by the use of tensor-product spaces to represent compound systems. As we noticed, such spaces allow for so-called 'entangled' (non-factorizable) states that imply strong (even perfect) correlations between remote systems. Again, the obvious ways to make sense of these correlations, in terms of a common causal history or processes that propagate purely locally in space-time, turn out not to work. Such is the upshot of a series of theorems, the first and most famous of which is due to Bell [1964]. The result is an apparent conflict with the tendency among space-time theorists of the past century or more to adopt some form of a principle of 'locality'. While various attempts have been made to resolve the apparent conflict, or to argue that it is not a problem, no very satisfying resolution seems

to be on offer.

2 WHENCE THE KINEMATICAL FORMALISM?

Our goal in this section is to review some attempts to say why we use the formalism that we do in quantum theory. There are numerous attempts to ‘derive’ the Hilbert-space formalism from physically ‘intuitive’ axioms. The attitude that we shall adopt here is that many of these attempts are worthwhile — they can contribute to our understanding of quantum theory — even if none of them succeeds in deriving the formalism from axioms whose physical import (much less truth) is always clear. It therefore makes sense to survey more than one such route to the formalism, rather than relying on just one of them. Here we will consider a few important representatives of the genre.

The first (§2.1) begins from the notion of a physical proposition, and argues that these propositions are properly represented by the subspaces of a Hilbert space. The second two routes (§2.2) begin from the notion of a physical state, and argue that these states are properly represented as probability measures over the subspaces of a Hilbert space. The final route (§2.3) is similar to these latter two, but while they attempt to characterize the state spaces in a somewhat abstract way (in terms of C^* -algebras), and then ‘represent’ them as states on a Hilbert space, this final route constructs the Bloch sphere (§1.3.5) directly, and builds (the state spaces of) higher-dimensional Hilbert spaces from it. (However, having said that, we will not in fact examine the proofs enough to notice these differences in any detail.)

Our intention is not to claim that any of these routes has, once and for all, illuminated the true reasons behind the use of Hilbert spaces in quantum theory. Rather, in each case, certain interesting or important aspects of the use of Hilbert spaces in quantum theory can be illuminated by following through some of the arguments that take us (ideally) from relatively uncontroversial and physically clear principles to the Hilbert space formalism. Our emphasis here will be on those aspects of each route that seem to be potentially helpful for understanding quantum theory, skipping over what appear to be merely technical conditions. In each case, our discussion will be necessarily brief (relative to the full story); proofs and technicalities will be omitted.

Finally, many concepts (such as the lattice of propositions, and the convex space of states) that are important for the foundations of quantum theory are discussed here. Readers who are not necessarily looking to learn about routes to the Hilbert space formalism will still learn something independently important from reading at least some parts of this section.

2.1 From Propositions to Hilbert Space

In this section, we trace a route to Hilbert space that begins with the logic of physical propositions. The goal, here, is to identify quantum theory (i.e., the

Hilbert space formalism) as, in a sense, the unique theory that satisfies certain logical constraints.²⁰

2.1.1 *The Lattice of Propositions*

2.1.1.1 Physical Meaning of Propositions We begin by taking as fundamental the set of propositions about a physical system at a time. The idea is that such propositions will take the form ‘the system has the property P at time t ’. (Hence, equivalently, one may take properties as fundamental. For each property, there is a corresponding proposition stating that the system possesses the property, at a given time. Here we shall speak solely in terms of these propositions.)

2.1.1.2 Definition It is standard in the algebraic approach to logic to assume that the set of all propositions (syntactically, sentences) forms a lattice (§7.4). The partial order in the lattice corresponds to implication: $P \leq Q$ means that P implies Q . The supremum (join) on the lattice corresponds to disjunction, and the infimum (meet) corresponds to conjunction. These identifications are far from arbitrary. Consider, for example, disjunction. The join of two elements, P and Q , in the lattice, L , is the logically weakest proposition implied by both P and Q , which is, arguably, just what one means by ‘disjunction’. Similar considerations apply to conjunction. Finally, assuming (again, as is standard) that the lattice is an ortholattice, the orthocomplement in the lattice corresponds to negation.

2.1.1.3 Motivation and Interpretation of the Lattice Operations Apart from the direct appeal to algebraic logic, such lattices have sometimes been motivated by an operational description of measurements. One speaks, for example, of yes-no experiments, to test for the truth of some proposition. The partial order on the set of such tests is derived from the idea that one test, Q , might be passed every time some other test, P , is, and in this case we write $P \leq Q$. However, I shall not pursue this approach in any detail here, but instead take it for granted that there is some way to understand these logical operations and relations in physical terms. (I do not mean to suggest that any particular approach — especially the various operational approaches — are entirely satisfactory. Indeed, one could argue that a firm understanding of the physical significance of the lattice-theoretic connectives is the Achilles heel of the route to Hilbert space via propositions.²¹)

²⁰There are many versions of this program, which got its start with Birkhoff and von Neumann [1936]. The discussion here is largely in the spirit of Piron’s [1976] work. A recent monograph covering much of the territory is that of Dalla Chiara, Giuntini, and Greechie [2004]. A more operational approach, somewhat different from these others, is that of Ludwig [1983].

²¹I do not mean to suggest, of course, that this issue has not been addressed by advocates of the quantum-logic program. Jauch [1968], for example, proposes an understanding of the conjunction of non-commuting projections (which, apparently, cannot be understood simply in terms of the proposition ‘both are measured’) in terms of a limit of repeated measurements.

2.1.1.4 Additional Constraints on the Lattice of Propositions In addition to the structure of the lattice of propositions that we have already established, we will need to assume a number of additional properties. First, we assume that the lattice has a bottom and a top (§7.4.5), 0 and 1, corresponding respectively to the logically false proposition and the logically true proposition. (We then require of the orthocomplement that $P \vee P^\perp = 1$ for all P . It follows that $0^\perp = 1$.)

Second, we assume that the lattice of propositions is complete and atomic (§7.4.5). This assumption amounts to the idea that there are some fundamental propositions — maximally specific propositions — and that one can interpret the truth of any weaker proposition as a consequence of the truth of some fundamental proposition. (It does not follow that one *must* interpret them in this way.) Another approach to justifying this assumption refers to the convex structure of the set of states. As we will see below (§2.2.1.3.c), there is good reason to suppose that this set has extremal points, corresponding to pure states, i.e., states of maximal information. That characterization suggests that each pure state should assign probability 1 to some maximal (logically strongest) proposition, which would be an atom in the lattice of propositions.

Next we must assume that the lattice is irreducible (§7.4.4). The assumption of irreducibility is far from trivial, but does have something approaching a physical interpretation, due to the following theorem from the theory of lattices:

An ortholattice is irreducible if and only if its center (§7.4.3) is trivial, i.e., $\{0, 1\}$.

Now, a proposition's being in the center means, logically, that it is compatible with all other propositions (§1.2.4). Therefore, one way to understand the irreducibility of a lattice is as a consequence of the assumption that its center is trivial, i.e., that *every* proposition (except 0 and 1) is incompatible with at least some other proposition. Note that this assumption is in a sense a 'maximal' violation of the classical law $p = (p \wedge q) \vee (p \wedge q^\perp)$, for it asserts that the law fails for *every* p (which is not to say that there are not *some* p, q for which the law holds, but that for every p , some q exists for which the law fails — recall, also, the discussion surrounding the conditions 23).

Our final assumption is that L satisfies the covering property (§7.4.5). The motivation here is, alas, less clear, although some arguments have been given.²²

2.1.2 Piron's Program

Piron's program was to characterize exactly those lattices that are (isomorphic to) the lattice of subspaces of a Hilbert space. He got as far as the following theorem:

Theorem (Piron [1964]) If L is a complete, atomic, irreducible, orthomodular lattice that satisfies the covering law and has at least 4 orthogonal atoms, then it is (isomorphic) to the lattice of subspaces of an inner product space, V .

²²For example: Piron [1964], and Cohen and Svetlichny [1987].

In fact, Piron was able also to say something about the field over which V is defined, but not much. In the end, Piron's theorem is suggestive, but far from the desired end, which is to characterize the lattice of subspaces of a Hilbert space as the unique structure satisfying certain logical constraints.

Piron's program made a huge step forward with the proof of the following theorem, due to Solér

Theorem (Solér [1995]) If the lattice L in the statement of Piron's Theorem contains an infinite orthonormal sequence, then the vector space in question is a Hilbert space over the reals, complex numbers, or quaternions.

This result is an important contribution to Piron's program, although there is an obvious limitation to its applicability: it does not cover the case of finite-dimensional Hilbert spaces (for example, those describing the spin of a particle, as in §1.3).

2.2 From States to Hilbert Space

2.2.1 An Approach in Terms of Convex Spaces of States

We will now consider two approaches to Hilbert space that begin with the structure of the space of states. The first begins with the observation that the states form a convex set (§2.2.1.1). Our procedure will be to embed an arbitrary convex set into a vector space, V (§2.2.1.2), and argue (in the subsequent sections) for progressively adding more structure to V , until it has enough structure to support a theorem to the effect that the convex set of states can in fact be represented as density operators on a Hilbert space.

2.2.1.1 Convex Spaces of States Earlier we noticed that given any two probability measures, p and p' , they can be combined to form a third, $q = rp + (1 - r)p'$, where $0 < r < 1$. It is easily checked that q thus defined satisfies Kolmogorov's axioms (§7.5.6), if p and p' do. Indeed, any convex combination (§7.1.10) of probability measures yields another probability measure, called a 'mixture' of the measures appearing in the convex combination.

The general idea, then, will be to take a convex space of states as fundamental. At this point, we make no commitment about the space of states other than that it is convex. We will then impose extra conditions on this space, eventually forcing our space of states to be, in fact, the space of states on a Hilbert space. In other words, we will have characterized the space of states (density operators) on a Hilbert space.²³

²³The discussion in this section is largely an expansion and explanation of the program outlined by Haag [1992, §VII.2]. The program has since taken on a new form, in the hands of Alfsen and Shultz [2003], as briefly mentioned below.

2.2.1.2 Embedding in a Vector Space So consider, to start, an arbitrary convex set, S . It is mathematically natural and convenient to embed S into a real vector space, V , in part because convex combinations are just a special type of real linear combination, and the latter is naturally defined in the context of a (real) vector space.

If S is generated by its extreme points, then our immediate task is easy. (An extreme point of a convex set is a point that is not itself a convex combination of other points in the set. S is generated by its extreme points just in case every element of S can be expressed as a convex combination of extreme points.) In that case, we can define the embedding by letting V be the free vector space generated by the extreme points of S . (Intuitively, V is then all formal real linear combinations of the extreme points of S .) However, this approach clearly requires that we assume that S has extreme points; and while it does in quantum theory as standardly understood, it is worth seeing that one need not assume so, but can instead derive this fact from other considerations. (In the context of a physical theory, where S is a set of states, the extreme points are just the pure states — recall, also, §2.1.1.4.) So let V be a real vector space having S as a subset (i.e., S is embedded in V — given the real linear structure of S , it should be clear that it is always possible to find a vector space that contains a subset that is isomorphic to S as a convex set; our problem here is not to show that there is such a V , but to construct it from S , as above).

2.2.1.3 Sufficient Conditions for the Existence of Extreme Points

2.2.1.3.a Observables We will take an arbitrary observable, f , of our theory to be a function from states to expectation values. That is, $f(v)$ is the expectation value of f in the state v . Indeed, after embedding the states in a vector space, V , we will consider the observables to be (possibly a subset of the) real-valued linear functionals on V . It is clear that observables, understood as maps from states to expectation values, should be linear functionals, because for any state, $v = \sum_n w_n v_n$ (written here as a convex combination of states v_n), we must have $f(v) = f(\sum_n w_n v_n) = \sum_n w_n f(v_n)$. Otherwise, $f(v)$ would in general be numerically different from the expectation value of f on a system randomly chosen from the mixture $\sum_n w_n v_n$.

Let O be the set of observables. We need not (yet) make any commitments about the contents of O except to require that it separate V ; that is, for any non-zero $v \in V$, there is an $f \in O$ such that $f(v) \neq 0$ (equivalently, if $v_1 \neq v_2$ then there exists $f \in O$ such that $f(v_1) \neq f(v_2)$).²⁴ In S , this condition amounts to requiring that O be rich enough to make probabilistic distinctions amongst elements of S . To extend to all of V , assume now that O separates S , and note that if O does not also separate all of V , then in fact V is ‘unnecessarily large’

²⁴This condition can always be met — i.e., one can always find some linear functional that does the job — for any normed vector space, by the Hahn-Banach theorem. Our space V will in fact be normed.

to embed S . For consider the subspace of all v such that $f(v) = 0$ for all $f \in O$. Then consider the quotient space V/W . There is a homomorphism from V to V/W that is an isomorphism from S to its image in V/W as convex sets. So we can just as well work with V/W , and doing so in fact eliminates structure from V that was not needed in order to embed S .

2.2.1.3.b Topology on the Embedding Space Therefore, we assume that O separates V . We now introduce a topology (§7.5) for V . Here, the guiding idea is that ‘infinitesimally small’ changes in the state should result in ‘infinitesimally small’ changes in expectation values.²⁵ We therefore introduce the coarsest topology on V that makes all of O continuous. Call it the O -topology.²⁶ In this topology, it makes a certain amount of sense to require that S be compact (§7.5.1).²⁷ For example, that assumption guarantees that (the expectation values of) observables are bounded, because the continuous image of a compact set is compact.

2.2.1.3.c The Existence of Extreme Points Given these assumptions about S and its embedding into a vector space, V , and the associated observables, O , we can apply the following theorem from functional analysis:

Theorem (Krein-Millman; see Rudin [1973, p. 70]) Let S be a convex subset of V , O a separating set of linear functionals on V , and let S be compact in the O -topology. Then S has extremal points, and it is the smallest closed convex set containing all of those points.

One says that S is the ‘closure of the convex hull of its extreme points’. In other words, S has extreme points, and is ‘generated’ by them in the sense that S is the closure of all convex combinations of extreme points. Hence every element of S can be written as a convex combination of extreme points (pure states), or is the limit of a sequence of such states.²⁸

2.2.1.4 Further Properties of the Embedding Recall the prescription for constructing the embedding space, V , as the free vector space over the extreme points of S (§2.2.1.2) Now that we have established that S indeed has extreme points, we can, almost, follow this prescription. The only addendum is that we

²⁵This assumption, while apparently natural, is certainly not compelling. See §2.3.2 for discussion of this idea in the context of Hardy’s approach.

²⁶A basis for this topology is given by all sets (‘open balls’), B , constructible as follows. Choose $v \in V$, $f_1, f_2, \dots, f_n \in O$, and $\epsilon > 0$. B is then the set of all $w \in V$ such that $|f_i(w) - f_i(v)| < \epsilon$ for $i = 1 \dots n$.

²⁷In the specific context in which we will soon work, compactness amounts to the assumption that the identity — the linear functional that has value 1 everywhere — is an observable.

²⁸Careful! Those familiar with quantum theory might be tempted to suppose that V is, or is isomorphic to, the Hilbert space containing the statevectors for a system. It is neither. As we will see (once we have defined a norm on V), it is a vector space in which the density operators form the unit sphere.

wish V to be closed under the O -topology. Hence, V is constructed as the closure of the free vector space over the extreme points of S . (By ‘closure’, here, we mean: ensure that the closure of every open set in V is also in V .) In this case, the pure states in S form a basis for V (because, in fact, the free vector space over the pure states is in essence ‘all formal linear combinations’ of pure states).

Let us now insist that O contain a linear functional, I , that assigns the value 1 to every pure state (extreme point) in S . This I is unique (because the pure states form a basis, and I is continuous by assumption and so is extended uniquely from the pure states to all of V). Now, let V^+ be the positive convex cone (§7.1.10) generated by the pure states. The states in V are then identified as the elements of V^+ that take the value 1 for I .

Indeed, I generates a norm on V in a natural way. For any $v \in V^+$, define the norm of v by $\|v\| = I(v)$. Now, any $v \in V$ can be decomposed as $v_1 - v_2$ for some $v_1, v_2 \in V^+$ (or is the limit of a sequence that can be so decomposed).²⁹ Hence we can define $\|v\|$ (now for any $v \in V$) as the infimum (over all such decompositions) of $\|v_1\| + \|v_2\|$, and require the norm to be continuous. V is thus a real Banach space (§7.1.7), and (by construction) the norm on V assigns norm 1 to every state. (We have now fulfilled our earlier promise that V would turn out to be a normed space.)

2.2.1.5 Faces and Propositions To what do these states assign probabilities? While we do not need to answer that question in order to finish the discussion (at this level), it is instructive to do so nonetheless: the ‘propositions’ to which states assign probabilities are faces of states in the convex set of states (§7.1.10).

In classical mechanics, one can form a face of a convex set by taking the closure, under convex combination, of a set of pure states. In quantum theory, the process of purification will in general add new pure states to the set, and so is essential for the construction of the face. (Recall, as well, that the classical states form a simplex, while the quantum states do not — see §1.2.7.) The physical idea behind the definition of a face in both cases is something like ‘the set of all states that one can create (from some initial set) via mixing, plus the set of all states of which the resulting mixed states could, in principle, be a mixture’.

There is a natural connection between faces of states and propositions about, or properties of, a physical system. (Here, as above, we shall use the terms ‘property’ and ‘proposition’ interchangeably. Also as before, we begin with the minimal assumption that the set of propositions is a poset, the partial order corresponding to implication.) In particular, let us say that a proposition (about a system, at a time) asserts that a measurement on the system reveals that the state of the system is in some face of states.

Note that an extremal point in the space of states is a (singleton) face, and vice versa. Hence one sort of proposition is of the form ‘the system is in the state v ’,

²⁹Sketch of a proof: the pure states span V , so write v as a linear combination of the pure states, and separate it into a part with positive coefficients and a part with negative coefficients. The former is clearly in V^+ . The latter, multiplied by -1 , is also in V^+ .

for any pure v . This association makes sense, because pure states are supposed to be, intuitively, states of maximal information, and if propositions correspond to faces, then the most specific propositions are the singleton faces (extremal points).

The idea behind the general association between propositions and faces is the following. Suppose you are handed an ensemble of systems, all in the same state, and asked to determine in what state the ensemble was prepared. Given the association between maximally specific propositions and pure states, if you can determine that some maximally specific proposition is true of every member of the ensemble, then you are done — the state is the corresponding pure state. But suppose that there is no such proposition, i.e., no maximally specific proposition true of *every* member of the ensemble. Then you have been handed a mixed state, but a mixture of what? In general, you will be able to determine with certainty only that the state was created by mixing states from the smallest face containing the actual mixed state (hence, the smallest proposition to which the actual state assigns probability 1).

In terms of standard quantum theory, the point here is that for any mixed state (density operator) W , the logically strongest proposition that is rendered certain by W is just $\text{ran } W$, here generalizing the notion of ‘range’ to mean ‘the image of W under all vectors in the Hilbert space’. (It should come as no surprise that $\text{ran } W = \bigvee_n P_n$, where P_n are the spectral projections of W .) In other words, $\text{ran } W$ is in fact the subspace of pure (vector) states that forms the smallest face containing W .

2.2.1.6 Hilbert Space representations of faces of cones

2.2.1.6.a Homogenous cones Consider the automorphisms on the set of states (that is, any map from the set of states to itself that preserves the convex structure). Such maps plausibly correspond to the possible state-transitions. Moreover, they are naturally extended to linear maps on V that are automorphisms of V^+ (i.e., such maps take V^+ to itself, and preserve convex structure of V^+ ; hence they preserve the fact that V^+ is the positive convex cone generated by the pure states). The extension to V^+ is effected simply by the condition that for any such automorphism, f , any real number r , and any state, v , $f(rv) = rf(v)$. Recalling that every element of V can be written as a linear combination of elements of V^+ , one can see that f is thus naturally extended to a linear transformation on V .

A cone is said to be ‘homogeneous’ with respect to this set, T , of transformations if for any two non-extremal points, v, v' inside the cone, there is a transformation in T that takes v to v' . The physical idea here is that there is *some* way for a system to evolve from any non-extremal point to any other non-extremal point.

2.2.1.6.b The Case of Finitely Many Pure States While we would not want to assume that there are only finitely many pure states, nonetheless the following theorem is extremely suggestive:

Theorem (Vinberg [1965]). The faces of a self-dual,³⁰ homogeneous cone with finitely many extremal points are in one-to-one correspondence with the subspaces of some Hilbert space (over the real numbers, complex numbers, or quaternions).

Hence, if one is convinced that the states of a theory must form a convex set with the properties discussed above, one is, in the sense given by this theorem, committed to the Hilbert space formalism.

2.2.1.6.c The General Case Alfsen and Shultz [2003, p. 414] have extended this program to eliminate the unrealistic (and indeed false) assumption that there are only finitely many pure states. Their main result involves a number of technical assumptions about the structure of the convex set of states that we do not have the space to articulate here. Moreover, in the end, they arrive not quite at a characterization of Hilbert space (at least not directly), but at a characterization of the state spaces of a C^* -algebra, although, via a GNS construction, they can get to Hilbert space from there.³¹

2.2.2 An Approach in Terms of Pure States and C^* -Algebras

In this sense (i.e., in dealing ultimately with the state-spaces of C^* -algebras), the theorem of Alfsen and Shultz is similar to a theorem due to Landsman [1998, Theorem 3.9.2 and Corollary 3.9.3], which we will now consider, albeit briefly and (again) in outline.

Landsman argues that the set of pure states should be endowed with two distinct structures, dynamical and probabilistic, and that the two must be connected in the right way. He then adds what is supposed to be the characteristically quantum condition (described below), and arrives at a theorem that characterizes the state spaces of C^* -algebras.

2.2.2.1 The Poissonian Structure

2.2.2.1.a State Spaces as Poisson Manifolds The first sort of structure corresponds to the dynamics that we discussed earlier. Recall our discussion of the Heisenberg equation (§1.5.2.3.c), and specifically the fact that (i times the) commutator has the algebraic form of a Poisson bracket. In outline, the general theory of dynamics generated by Poisson brackets goes as follows.

We begin with some space (indeed, manifold — §7.5.2) of states, M . Given M , one defines the evolution of observables in terms of a Poisson bracket defined on the infinitely-differentiable real-valued functions on M , $C^\infty(M)$ (the observables), as a bilinear operation, $\{, \}$, on $C^\infty(M)$. (Recall (64), and see Butterfield, Ch. 1,

³⁰We pass over the issue of the physical motivation for this condition. See §7.1.10.

³¹For more on C^* algebras, GNS-constructions, and related matters, see Landsman, Ch. 5, Halvorson, Ch. 8, and Emch, Ch. 10, this volume.

§5, this volume, for further discussion.) Together with a choice of a scalar function $H : M \rightarrow \mathbb{R}$ as the Hamiltonian, one can use this definition of the evolution of observables to define an evolution (of states) on M (analogously to the equivalence between the Heisenberg and Schrödinger pictures). For a given Hamiltonian, the possible dynamical paths from M , thus defined, are the ‘Hamiltonian curves’ of the system.

Keep in mind that this construction is supposed to be very generic. We are essentially laying down ‘what one means’ by a certain kind of dynamical system. To this end, Landsman puts additional constraints of ‘well-behavedness’ on the space of states as a Poisson manifold, which we bypass here.

2.2.2.1.b Symplectic Leaves It can happen that some portions of a Poisson manifold are ‘inaccessible’ from others. In particular, there may be no (piecewise smooth) Hamiltonian curve connecting them, for any Hamiltonian. Let us say that two points, $x, y \in M$ of a Poisson manifold $\langle M, \{, \} \rangle$ are *symplectically equivalent* if for some Hamiltonian they lie on a single (piecewise smooth) Hamiltonian curve. This relation is clearly an equivalence relation, and therefore partitions M into ‘symplectic leaves’ (see Butterfield, Ch. 1, §5.3.3, this volume).

2.2.2.2 The Transition Probability Structure

2.2.2.2.a State Spaces as Transition Probability Spaces The only dynamical structure on a classical Hamiltonian phase space is its Poissonian structure (again see Butterfield, Ch. 1, especially §5.2.4, this volume). However, in standard quantum theory, there are two types of evolution: the continuous, deterministic, evolution described by the Poisson bracket, and the discontinuous, stochastic, ‘quantum jumping’ from one state to another, often associated with ‘measurement’. (See §5.4.3.)

This latter structure is reflected in the fact that, in Landsman’s scheme, the quantum state space must be a ‘transition probability space’, meaning that there must be maps, p , from pairs of elements in the space to $[0, 1]$ that satisfy: $p(v, w) = 1$ if and only if $v = w$; and $p(v, w) = 0$ if and only if $p(w, v) = 0$. (The expression $p(v, w)$ is read as ‘the probability of a transition from v to w ’.) In addition, we require that these probabilities be symmetric: $p(v, w) = p(w, v)$.

2.2.2.2.b Sectors A sector of a transition probability space is a region of the space that is isolated from the rest. That is, for Q a sector of states, $p(v, w) = 0$ for all $v \in Q$ and all $w \in Q'$ (the complement of Q in the set, S , of all states). Note that the symmetry of p implies that a system cannot make a transition to or from a sector.

2.2.2.2.c Superpositions Let $Q \in S$. We define Q^\perp by

$$(65) \quad Q^\perp := \{v \in S \mid p(v, w) = 0 \forall w \in Q\}.$$

That is, Q^\perp is the set of all states that are inaccessible (via a single probabilistic transition) from every state in Q . We can use this definition to characterize a generic notion of a ‘superposition’, as follows: the class of all ‘superpositions’ of the states v and w is $\{v, w\}^{\perp\perp}$ (see [Butterfield, 1993]).

2.2.2.3 Landsman’s Theorem Landsman shows, roughly, that the state space of a C^* algebra is uniquely determined by its Poissonian and transition probability structure, assuming (among other things) that symplectic leaves correspond to sectors. He then characterizes *quantum* theories as those in which the ‘2-sphere’ property holds, i.e., condition (iv) in the theorem as stated below. *Classical* theories, on the other hand, are characterized by the condition that the transition probabilities are trivial, i.e., $p(w, v) = \delta_{vw}$. Notice that in this case, the sectors are singletons. We will see a similar characterization of the difference between quantum and classical systems below (§2.3.2).

In effect, then, Landsman [1998, 104–106] proves roughly the following theorem:³²

Theorem (Landsman [1998, 104–106]). A pure state space, S , is the pure state space of a quantum system if and only if: (i) S is a Poisson manifold; (ii) S is a transition probability space; (iii) the symplectic leaves of S correspond to the sectors of S ; (iv) for any $v, w \in S$, $\{v, w\}^{\perp\perp}$ is isomorphic as a transition probability space to the space of statevectors in \mathbb{C}^2 .

I have already discussed conditions (i) and (ii). Condition (iii) is a requirement that what cannot happen by continuous evolution cannot happen by stochastic evolution either (and vice versa). In other words, if it is ‘dynamically impossible’ to get from state v to state w , then the probability of a stochastic transition from v to w is 0 (and vice versa).

Condition (iv) is, of course, the one that does a lot of the work getting us to the Hilbert space formalism, for it is essentially the requirement that the set of all ‘superpositions’ of a pair of states forms a transition probability space that looks like the quantum-mechanical pure states on the space \mathbb{C}^2 , discussed earlier (§1.3.3.2) in the context of spin. Whether or to what extent this more or less explicit reliance on quantum theory is ultimately satisfactory is a matter of taste, but it is, at any rate, worth noticing that what is ‘essentially quantum’ about the Hilbert space formalism can (more or less) be reduced to this assumption.

2.3 Hardy’s Axioms

Finally, we consider a result from Hardy [2001; 2002], again in the tradition of trying to explain the origins of the quantum formalism. This approach also begins

³²The conditions of the theorem are not stated here in full rigor. In addition, the proof of the theorem requires a few other technical assumptions whose immediate physical import is perhaps not clear. I have left them out.

from the notion of a state space, but its framework differs sufficiently from the previous two approaches that we consider it separately.

2.3.1 The Framework

2.3.1.1 States as Probabilistic Predictors Much as has been done already in the previous sections, Hardy takes *states* in a physical theory to be the determiners of probabilities associated with each possible outcome of any measurement that may be performed by a given preparation of the system. (States are thus associated with preparations.) Hence, for a given preparation, knowing the associated state allows one to predict the probabilities for the results of any measurement.

2.3.1.2 Degrees of Freedom One mathematical characterization of a state is as a ‘list’ of all of these probabilities. (Of course, in general there are at least uncountably many items in the ‘list’.) However, in general the state space has some structure that allows states to be somehow characterized by a shorter list of what Hardy calls ‘fiduciary’ probabilities. In a given theory, we define the *degrees of freedom*, K , to be the smallest number of fiduciary probabilities that is sufficient to determine the state.

2.3.1.3 Dimension In addition, there may be sets of states that can be distinguished one from another with probability 1 in a single measurement. In other words, for each pair of states, v, w , in the set, if v assigns non-zero probability to some outcome of the measurement, then w assigns probability zero to the same outcome. There will in general be a maximum number, N , of states that are distinguishable in this way. Hardy calls N the *dimension* of the space.

2.3.2 The Axioms

Hardy proposes five ‘axioms’. The first axiom simply underwrites our earlier assumption that states can be associated with preparations, that is, that there are stable relative frequencies for the outcomes of measurements for a given type of preparation. The remaining axioms are, as adapted from Hardy (2001):

Subspaces For any integer, N , there exist systems with dimension N . Moreover, all systems of dimension N , and all systems with higher dimension but whose state restricts the system to an ‘ N -dimensional subspace’, have the same properties.

Composite Systems A composite system consisting of systems A and B , with degrees of freedom and dimensions K_A, K_B and N_A, N_B respectively, have $K = K_A K_B$ degrees of freedom and dimension $N = N_A N_B$.

Continuity For any dimension, N , any system of dimension N , and any two pure states, v and w , of such a system, there exists a continuous reversible transformation (continuously connected to the identity) from v to w .

Simplicity For given N , K takes the minimum value consistent with the other axioms.

The motivation for some of these axioms is relatively clear, and for others, not as clear. We will not discuss them all here. ‘Simplicity’ comes into play because the other axioms imply that $K = N^m$ for some integer m . For $m = 1$, the continuity axiom is violated, and one arrives at classical probability theory. For $m = 2$, one arrives at quantum theory.

The continuity condition is clearly significant, therefore. One understanding of it that Hardy has encouraged is that it expresses the desideratum that ‘small changes’ in the state should entail ‘small changes’ in the predictions based on that state. However, it is not clear whether this principle is physically compelling. After all, in Hardy’s sense, the state space of classical physics is *not* continuous (cf. the beginning of §2.2.2.3), and yet in general one does not suppose that there is somehow a serious problem with the relationship between changes of state and changes of prediction based on that state.

A different understanding of continuity connects it with superpositions. The basic point is easiest to visualize geometrically, and we will leave the matter at that here. The quantum state space is ‘continuous’ (in Hardy’s sense) because for any two pure states, there is another pure state that is ‘between’ them, and in fact this ‘middle’ state is a superposition of the two original states. In other words, continuity holds precisely because the superposition principle holds. Continuity fails in the classical theory because the superposition principle fails there. From this point of view, it is less surprising — though not necessarily less important — that continuity is what makes the difference, in Hardy’s framework, between classical and quantum theories.

3 EMPIRICAL CONTENT

In this section, we will discuss the issue of how the formalism of quantum theory gets empirical content. That discussion will lead naturally to a discussion (§4) of uncertainty, because the particular way that we will, here, understand how the formalism gets empirical content leads naturally to the uncertainty principle (§4), as it should. Hence this section is not merely expository. Implicitly, we will develop an argument in favor of a particular way of understanding how the formalism gets its empirical content, based on the fact that it leads naturally to uncertainty.

One question about empirical content concerns how measurement (or empirical observation more generally) is modeled in the theory. Without such a model, it is hard to see how the theory can make predictions about empirical observations. Our first task (§3.1) will be to say something about this issue. But the more difficult — and arguably more philosophically interesting — issue concerns how *any* element of the formalism is connected with empirical fact at all. The remainder (and majority) of this section will concern that issue.³³

³³Much of that material is adapted and revised from Dickson [2004a; 2004b].

There are two elements in the present approach to establishing a connection between the formalism and its empirical content: symmetries, and reference frames. These elements are connected — ‘legitimate’ reference frames are connected to one another via certain symmetry transformations — but I will, for the most part, keep them separated. In this section, I first frame the issue of empirical content in terms of POVMs (§3.2), then discuss the role of symmetries (§3.3) and reference frames (§3.4) in the definition of physical quantities, and finally I briefly sketch an account (§3.5) of how the formalism gets its empirical content.

3.1 Measurement

3.1.1 The Standard Account of Measurement

One frequently encounters the following account of measurement in quantum theory. Suppose that the state of the system is W (a density operator). Suppose that one measures an observable represented by the POVM $E : \mathcal{B}(S) \rightarrow \mathcal{B}(\mathcal{H})^+$. Then the result of the measurement will be some $\Delta \subset S$, with probability $\text{Tr}[E(\Delta)W]$. Furthermore, the state after measurement is just the ‘projection’ of the original state onto the result. In the case where the POVM is a PVM, the projection is given by $E(\Delta)WE(\Delta)$ (ignoring normalization). Otherwise, the usual procedure is to define ‘measurement’ operators $M(\Delta) = \sqrt{E(\Delta)}$, in terms of which the ‘projection’ (a misnomer, in this case) is (again ignoring normalization) $M(\Delta)WM^*(\Delta)$. Notice that the latter prescription is equivalent to the former for PVMs.

While useful for making predictions, this account of measurement is completely unsatisfactory, from a foundational point of view. One problem, as it has been pointed out repeatedly by both physicists and philosophers, is that measurement is itself a physical process, and in particular a physical interaction between two (or more) physical systems, one of which we call an ‘apparatus’ and the other of which we call the ‘measured system’. Hence, if quantum theory is our best theory of interactions amongst physical systems, it should be capable of describing this interaction in such a way that the result — the states of the two systems at the end of the interaction — are as stated above.

3.1.2 Impulsive Measurement

There is indeed an account of measurement along those lines.³⁴ It is the so-called ‘impulsive model’ of measurement. Suppose we are going to measure an observable represented by the self-adjoint operator, F . For simplicity, suppose that the measured system is initially in the pure state $|\psi\rangle$ and the apparatus is in the pure state $|\chi\rangle$. The compound system is thus in the state $|\Psi\rangle = |\chi\rangle|\psi\rangle$. Suppose that apart from the interaction between them, each of these systems evolves freely, with (free) Hamiltonians H_S and H_M . The total Hamiltonian for the combined system is thus $H_{\text{total}} = H_S + H_M + H_I$, where H_I is the interaction Hamiltonian

³⁴For an extended classic discussion, which is more or less followed here, see Bohm [1951, ch. 22, especially §5].

(i.e., it represents the energy exchanged between the systems). Finally, let Π be the momentum observable for the apparatus.

Now, for our model of measurement, we take $H_I = g(t)\Pi \otimes F$ (henceforth the \otimes is left implicit), where $g(t)$ is an interaction function given by $g(t) = \gamma f(t)$, with γ a constant, the ‘interaction strength’, and $f(t)$ is a function that is zero except for t between 0 (when the measurement–interaction begins) and τ (when the measurement ends), and $f(t)$ is (purely for convenience) normalized, i.e., $\int_0^\tau f(t)dt = 1$. The Schrödinger equation (1.5.2.3.a) may therefore be written:

$$(66) \quad \frac{d}{dt}|\Psi\rangle = -i\left(H_S + H_M + \gamma f(t)\Pi F\right)|\chi(t)\rangle|\psi(t)\rangle,$$

where now we have explicitly indicated the time-dependence of the states of the measured system and the measuring device.

For an ‘impulsive’ measurement, τ is very small and γ is very large. (The interaction is quick and strong.) If we may assume that the apparatus and system have low or zero momentum, then during the interval $[0, \tau]$, the influence of the interaction Hamiltonian on the evolution of the compound system completely swamps the influence of the free Hamiltonians, so that, during this period, we have

$$(67) \quad \frac{d}{dt}|\Psi\rangle \simeq -i\gamma f(t)|\chi(t)\rangle|\psi(t)\rangle.$$

We can readily solve this equation³⁵ to get the compound state at the end of the interaction (*immediately* at the end, before the free Hamiltonians take over again):

$$(68) \quad |\Psi(\tau)\rangle \simeq \exp\left[-i\int_0^\tau \gamma f(t)\Pi F dt\right] |\chi(0)\rangle|\psi(0)\rangle.$$

Write $|\psi(0)\rangle$ in terms of the (normalized) eigenstates, $|f_n\rangle$, of F (which, for simplicity, we will assume to be maximal): $|\psi(0)\rangle = \sum_n \langle f_n|\psi(0)\rangle|f_n\rangle$. (See the last paragraph of 7.1.4). Then (68) becomes

$$(69) \quad |\Psi(\tau)\rangle \simeq \sum_n \langle f_n|\psi(0)\rangle \exp(-i\gamma f_n\Pi)|\chi(0)\rangle|\psi(0)\rangle,$$

where f_n is the eigenvalue of F corresponding to the eigenvector $|f_n\rangle$. Now define

$$(70) \quad |\xi_n\rangle = \exp(-i\gamma f_n\Pi)|\chi(0)\rangle.$$

Because γ is large, these states are effectively orthonormal:³⁶

$$(71) \quad \langle \xi_n|\xi_m\rangle = \langle \chi(0)| \exp[-i\gamma(f_n - f_m)P]|\chi(0)\rangle \simeq \delta_{nm}.$$

³⁵The solution in (68) follows only because the interaction Hamiltonian is given in terms of a scalar function of time. General time-dependent Hamiltonians cannot be treated in this way. See, e.g., Cohen-Tannoudji [1977, 172–175].

³⁶The second equality in (71) holds by the Riemann-Lebesque lemma, under the assumption that the apparatus has a continuous spectrum for P .

As we will see below in a more generic context (78), because Π is the ‘momentum’ of a pointer, the states $|\xi_n\rangle$ are ‘spatial translations’ (i.e., translations in the value of the ‘pointer-position observable’) by an amount γf_n (and therefore, for large γ , they are macroscopically distinguishable states). The upshot of this discussion is that the final state of the compound system is

$$(72) \quad |\Psi(\tau)\rangle \simeq \sum_n \langle f_n | \psi(0) \rangle |f_n\rangle |\xi_n\rangle.$$

Notice that this state is entangled, and indeed that it represents a perfect correlation (§1.2.6.4) between the value of F for the measured system and the pointer-position.

In order for this model to match the prescription of the previous subsection, we must adopt from that prescription the ‘rule’ that when the apparatus shows the result $|\xi_n\rangle$, the state of the system ‘collapses’ to $|f_n\rangle$. (Or, the state of the compound system is projected onto $\mathbb{I} \otimes |\xi_n\rangle\langle\xi_n|$, which, because of the perfect correlation, will have the effect of making the compound state into the product state $|f_n\rangle|\xi_n\rangle$.) I will discuss this ‘collapse’ rule in §5.4.3.

Does this model apply to all, or indeed to any, actual ‘measurements’ made by actual physical devices? To a reasonable degree of accuracy, probably. But the point here is not to make such a claim, and certainly not to catalogue the varieties of physical measurement and the ways in which they might be modeled quantum mechanically. Creating and justifying such models is the business of physics. The point, here, is that quantum theory can, to some degree, supply a model of measurement. It need not rely on the non-account of the previous subsection.

One final observation about this model: it models a type of measurement in which the state of the system is in a sense unchanged by the interaction with the measuring device: if the measurement is repeated, the probability of getting the same result the second (and subsequent) times is 1. Notice that, contrary to the occasional declaration by physicists or philosophers, in this sense at least, measurement does *not* necessarily ‘disturb’ the state of the measured system. Following Pauli [1958], such measurements are often called ‘measurements of the first kind’.

Of course, measurement sometimes does disturb the measured system. Indeed, sometimes it destroys the measured system. Measurements in which the state of the system is disturbed by the process of measurement (i.e., measurements that either are not repeatable, or whose results will not necessarily be the same upon repetition) are often called (again following Pauli) ‘measurements of the second kind’.

3.1.3 Weak Measurement

Once we begin modeling measurement as an actual physical process, it becomes natural to ask what would happen if the physical circumstances were different. One natural case to consider is where the measurement is ‘adiabatic’; that is, the interaction is weak, and takes a long time (on some appropriate scale).

One scheme for realizing this idea has been called ‘protective measurement’ [Aharonov *et al.*, 1993]. The idea is nicely illustrated by the case where $|\psi(0)\rangle$ is the ground state of a harmonic oscillator (the crucial feature of which, for us, is that there is a finite energy difference between the possible states). When τ is large and γ is small, of course we cannot ignore the evolution due to the free Hamiltonians. Let the interaction Hamiltonian be $H_I = g(t)QF$. The solution to (66) in this case is

$$(73) \quad |\Psi(\tau)\rangle = \exp[-i(H_S\tau + H_M\tau + \int_0^\tau \gamma f(t')QF dt')] |\chi(0)\rangle|\psi(0)\rangle$$

$$(74) \quad = \sum_n \exp[-i(H_S\tau + H_M\tau + \gamma QF)] \langle f_n | \psi(0) \rangle |\chi(0)\rangle |\psi(0)\rangle.$$

Now, because there is a finite difference in energy between the ground state and any excited state of the measured system, one must add a finite amount of energy to the ground state in order to change it. According to the quantum-mechanical adiabatic theorem (Schiff [1968, 289–291]), if the energy added to a system is small enough and spread out over a long enough time, it is not additive (i.e., the total energy added does not get larger and larger), but adiabatically negligible. Indeed, the probability amplitudes for states other than the ground state can be made arbitrarily small, with suitably small γ and large τ . In other words, the term $\exp[-i\gamma QF]$ has no net effect on $|\psi(0)\rangle$, and we need consider its effect only on $|\chi\rangle$. (Note that there *is* a possible effect on $|\chi\rangle$ if we presume that the energy spectrum for $|\chi\rangle$ is continuous, or effectively so.) Hence (74) becomes

$$(75) \quad |\Psi(\tau)\rangle \simeq \sum_n |\phi_n(\tau)\rangle |\psi(\tau)\rangle := |\chi_n(\tau)\rangle |\psi(\tau)\rangle,$$

where we have defined $|\phi_n(\tau)\rangle = \exp[-i(\gamma f_n Q + H_M\tau)] |\chi(0)\rangle$.

Notice that the state (75) is a product state — the interaction effects only the apparatus, and does not entangle the measured system with it. To see how this change in the state of the apparatus can be used to gain information about the system, recall (59). Taking the expectation value of $\mathbb{I} \otimes \Pi$ (where here Π is the momentum conjugate to Q) in the state $|\Psi(\tau)\rangle$ from (75), we find that

$$(76) \quad \frac{d}{dt} \langle \chi(\tau) | \mathbb{I} \otimes \Pi | \chi(\tau) \rangle = -g(t) \langle \psi(\tau) | F \otimes \mathbb{I} | \psi(\tau) \rangle.$$

In other words, the expected value of momentum for the apparatus is an ‘indicator’ of the expected value of F for the measured system. If the system is in a stationary state (i.e., $|\psi(t)\rangle = |\psi(0)\rangle$), then we could, for example, make many ‘protective’ measurements of F , measure the average value of the momentum of the apparatus after the interaction, and gain information about the expected value of F for the measured system.

There are two important remarks to make about this scenario. First, as many have pointed out, in order for the scheme to work, one must know the state of the measured system in advance. (In particular, in this case one must know that it is the ground state of the harmonic oscillator.) Otherwise, we will not know that it is ‘protected’ (i.e., will not change state as a result of the interaction). So there

is an important sense in which the protective measurements do not yield any new information about the system. Moreover, it has been argued [Uffink, 1999a] that only observables, F , that commute with the measured system's Hamiltonian can be measured in this way.

Nonetheless, despite these limitations, there remains the question of what exactly is going on in such an interaction. In the case of impulsive measurements (of the first kind, anyway), if we measure F on a system whose value for F was just measured, we will gain no new information about the value of F for the measured system. Nonetheless, we might be tempted to explain the result of the second measurement thus: the measuring device interacted with the system in such a way that its indicator-state became correlated with the state of the system, and in particular with its value for F . That is, the second interaction was, again, physically a measurement, albeit one that was bound to give us no new information. One might, then, be tempted to say a similar thing about protective measurements. True, they tell us nothing that we did not already know. But how are we to understand what is going on during such interactions? One plausible understanding is that the apparatus' state is changing because it is somehow sensitive to the expected value of F on the measured system (even though, of course, we already *know* what that expected value is). Indeed, what else *could* explain the change in the apparatus' state?

Second, protective measurements are an explicit model of a more general class of measurements, so-called 'weak measurements'.³⁷ The general scheme (somewhat in parallel with the general scheme for standard measurements, of which impulsive measurement is one model) is as follows. Consider a quantum system known to be in the state $|\psi_1\rangle$ at time 0, and known to be in the state $|\psi_2\rangle$ at time t . Typically, this knowledge is obtained by what is often called 'pre- and post-selection'. That is, the state $|\psi_1\rangle$ is 'pre-selected' (prior to what will be the 'weak measurement' of F on the system) by performing a standard (e.g., impulsive) measurement on some ensemble of systems and selecting just those for which the result of this first measurement is $|\psi_1\rangle$. (In particular, one might measure the observable corresponding to $|\psi_1\rangle\langle\psi_1|$ then select just those systems for which the result is 1.) Then, *after* the 'weak measurement', one again performs a standard measurement on the ensemble, selecting just those systems for which the result corresponds to $|\psi_2\rangle$. The resulting ensemble of systems is said to be 'pre- and post-selected'. For any given observable represented by the operator F , define the 'weak value of F on the pre- and post-selected ensemble' by

$$(77) \quad \frac{\langle\psi_2|F|\psi_1\rangle}{\langle\psi_2|\psi_1\rangle}.$$

Note that in the protective measurement discussed above, we presumed that the state of the measured system was unchanged over time, so that the weak value of F was just its expectation value in the state $|\psi(0)\rangle$.³⁸

³⁷ An early paper on weak measurements is [Aharonov *et al.*, 1987]. See also note 38.

³⁸ The very notion of 'weak values', not to mention the interpretation of these values, is con-

There is, of course, a great deal more to be said about measurement. Later (§5), we will consider perhaps the most important philosophical issue concerning measurement, namely, the measurement problem (which has already made an early appearance in the form of our barely suppressed skepticism about the collapse of the state after an impulsive measurement). For now, however, we will rest content with the observation that, barring problems to arise later, quantum theory provides a rich framework in which to describe measurements.

3.2 *The Issue of Empirical Content in Terms of POVMs*

Describing measurements is part, but only part, of the story about how the formalism is connected with empirical observation. Another part of the story concerns a more general question about the connection between formalism and physical fact. For example, we have been allowing observables such as S_u to ‘represent’ spin in the u -direction, but what precisely *is* this relationship of ‘representation’? How may the connection between formalism and physical fact be made, or understood? (Notice that the accounts of measurement above already presuppose an answer to this, more fundamental, question.)

It is crucial to understand that the issue here is not about how to engineer a spin-measuring device, for example. Rather, it is about what it *means* to ‘have’ spin-up in the u -direction (for example), and how this meaning is captured in the formalism. Supposing that there are no limitations of engineering, there remains a question about what laboratory procedures correspond to ‘measuring S_u ’, for example. Below, we will offer a partial answer to this question.

Recall that, considered as POVMs, observables are maps from (Borel sets of) ‘the possible values’ to positive operators. Another way to put the question above, then, is in terms of the empirical meanings of the mathematical elements in the domain of this map. (Once this question is answered, then, for example, the probability calculus associated with the elements in the range of the map becomes a calculus with empirical content.) Indeed, one advantage of conceiving of observables as POVMs (apart from the greater generality of this approach) is that it affords greater precision to a discussion of the issue of which parts of the formalism are, in which contexts, doing the work of ‘representation’. For a POVM $E : \mathcal{B}(S) \rightarrow \mathcal{B}(\mathcal{H})^+$, the elements of $\mathcal{B}(\mathcal{H})^+$ are doing the representing, and the elements of $\mathcal{B}(S)$ (indeed, ultimately, S) are, in a sense, what is being represented.

Why we represent with elements of $\mathcal{B}(\mathcal{H})^+$ is a question that we have addressed already in §2 (though certainly it has not been conclusively answered). But what is the physical interpretation of the domain of the map E , and how are we to

troversial. A recent discussion by the main proponent (and co-author), with plenty of references to prior work, both pro and con, is Aharonov and Botero [2005]. Note that Aharonov and other proponents often discuss weak values in the context of a ‘two-state-vector’ formalism for quantum theory (which concerns pre- and post-selected systems and is supposed to be time-reversal invariant) that is itself controversial; however, the notion of a weak value is not irrevocably tied to that approach to quantum theory, but only to the idea of a pre- and post-selected ensemble, the operational meaning of which, at any rate, is clear enough.

understand the statement that some element in its domain is ‘represented by’ some element in its range? What, in other words, is the relationship between the elements of the mathematical formalism that we have described and physical matters of fact? And finally, why do we pick one map (POVM) rather than another to represent some given physical quantity? The next section (§3.3), on symmetries, and the following one (§3.4), on reference frames, introduce material that will eventually contribute to a sketch of a story (§3.5) about the empirical content of quantum theory, one that addresses, or suggests ways to address, some of the issues raised in this section.

3.3 *Symmetries*

3.3.1 *Groups as Symmetries*

There is a traditional account of one way that groups have been related to empirical content.³⁹ Take any group, G , and consider its action on a set, S . If two elements of S are connected by an element of G , then call them ‘equivalent’. One can readily verify that G thus partitions S into equivalence classes, and we can say, then, that G is a group of ‘symmetries’ on S , in the sense that elements of S connected by an element of G are in some important sense ‘the same’. (If, according to some theory of the elements of S , distinct elements within the same equivalence class in S can have ‘importantly’ distinct properties, then in this theory, G would arguably *not* be a symmetry. What one means by ‘important’ is subtle, but in the context of this discussion it could, for example, mean ‘empirical’, or ‘observable’.) For example, let the elements of S represent the positions of all of the particles in the universe (i.e., S is a configuration space for the universe). A spatial translation of a point in S results in a universe that is, arguably, no different, empirically (because all distances and other spatial relations amongst the particles stay the same).

3.3.2 *Groups and Observables*

The (outline of an) approach to empirical content that we shall propose is given, in part, in terms of groups of transformations on S (considered as the domain of a POVM, E), and the requirement that E in a sense preserve the behavior of S under those transformations.⁴⁰ To get this view off of the ground, we require (at least!) the following two things of a quantum theory: (1) a faithful representation of the relevant groups on a Hilbert space, and (2) a map from $\mathcal{B}(S)$ to $\mathcal{B}(\mathcal{H})^+$ that in some relevant sense ‘preserves’ the action of these groups.

³⁹Here I am considering groups primarily as transformations on the set of states. One can also think of them in terms of formal transformations of physical laws. For more discussion of this and related points, see Brading and Castellani, Ch. 13, this volume.

⁴⁰The view described here has been strongly influenced by discussions with Scott Tanona, and by Tanona [2002; 2006]. (Tanona’s view is distinctive, and differs from the one presented here in important ways.) Indeed, it is largely due to those discussions that I began to develop a view about this matter.

Let us consider these points in the context of an example, the Galilean group, \mathcal{G} , and more specifically, the spatial translations, \mathcal{A} , and boosts, \mathcal{V} (i.e., ignoring rotations and temporal translations). For any $a \in \mathcal{A}$ and $b \in \mathcal{V}$, let U_a and V_b be the corresponding elements of a faithful representation of \mathcal{A} and \mathcal{V} , respectively. It is, at this point, an open question whether there *exist* faithful representations of \mathcal{A} and \mathcal{V} on a given Hilbert space, i.e., whether requirement (1), above, can be met. On an infinite-dimensional space, there are in fact operators P (momentum) and Q (position) that do the job:

$$(78) \quad \begin{aligned} U_a &= e^{-iPa} \\ V_b &= e^{-iQmb}, \end{aligned}$$

respectively, where m is the mass of the particle, and appears here because momentum is mass times velocity.⁴¹

Notice that positions are translated by the action of \mathcal{A} and invariant under the action of \mathcal{V} . The converse holds for velocities (and therefore, of course, momenta).

Now, let $E_Q : \mathcal{B}(\mathbb{R}^3) \rightarrow \mathcal{B}(\mathcal{H})^+$ be the POVM for position and consider the action (§7.6.7) of \mathcal{G} on \mathbb{R} (or \mathbb{R}^3 — see note 41), understood as representing *positions* of a particle.⁴² In that case, the requirement (2) above amounts to the requirement that the quantum-theoretic representation of position have the same properties:

$$(79) \quad \begin{aligned} U_a E_Q(\Delta) U_a^{-1} &= E_Q(\alpha_a(\Delta)) \quad (\text{covariance}) \\ V_b E_Q(\Delta) V_b^{-1} &= E_Q(\Delta) \quad (\text{invariance}) \end{aligned}$$

for any $a \in \mathcal{A}$ and any $b \in \mathcal{V}$, where α_a is the action of a on $\mathcal{B}(\mathbb{R}^3)$. Corresponding conditions must hold for E_P , the POVM for momentum (i.e., it should be invariant under translations and covariant under boosts). Note that at this point, we do *not* presume that the P and Q in (79) are those from (78). Indeed, these conditions (79) turn out to be powerful enough on their own to determine the maps E_Q and E_P , as well as the representations (78) (up to unitary equivalence). We shall return to this point below. Here, we take it as motivation for *defining* ‘position’ and ‘momentum’ in terms of (79) (and the corresponding conditions for momentum).

One might have either of two objections at this point. First, where did these requirements come from? Why do we insist that position must have this particular

⁴¹ One normally says that position is the ‘generator of translations’ and momentum is the ‘generator of boosts’. The reason ultimately has to do with the fact that the momentum (position) operator is involved in the expression for an infinitesimal translation (boost). The expressions in the text for finite translations and boosts are essentially integrals of their infinitesimal counterparts. Note, also, that (78) is given in one spatial dimension only. Replacing P and Q with the ‘vectors’ of operators $\vec{P} = (P_x, P_y, P_z)$ and $\vec{P} = (P_x, P_y, P_z)$, and the parameters a and v with vectors (from \mathbb{R}^3) \vec{a} and \vec{b} , we would have the three-dimensional version. Note, for example, that these changes do not change anything about the definition of a continuously parametrized group in §7.6.4.

⁴² \mathcal{G} naturally acts on the phase space in classical physics, but we can of course consider its action on the reduced space of position or velocity.

invariance and covariance, even in the classical case, or at all? The answer is that, at least in this discussion, we are taking covariance with respect to \mathcal{A} and invariance with respect to \mathcal{V} to *define what we mean* by position. For example, if the position of a particle is \vec{x} in a reference frame F (below we will consider in more detail the role of reference frames in this discussion), and if frame F is related to frame F' by a spatial translation \vec{a} , then the position of the particle in frame F' is $\vec{x} - \vec{a}$. The seeming triviality of this feature of position is, one might suppose, a consequence of the fact that it is part of what ‘position’ *means*. That is, an observable (POVM) that lacked this feature would, *ipso facto*, not be ‘position’. Similar remarks will hold for boosts, and again for the relationship between momentum on the one hand and translations and boosts on the other. Corresponding remarks hold for other observables, such as angular momentum and spin.

Second (objection), by insisting that position ‘mean the same thing’ in the context of quantum theory (i.e., that the POVM respect, in the relevant sense, the action of the (representation of the) Galilean group), are we not thereby preventing ourselves from learning, perhaps, that position is ‘very different’ from what we thought it was? (A similar remark holds, of course, for momentum, and indeed for the hosts of other physical quantities that can be defined in this way.) There are two answers to this objection. First, in the context of quantum theory (where the position and momentum POVMs *do* obey the ‘correct’ invariances and covariances), we *do* in fact learn that position is ‘very different’ from what we thought it was, and this fact already makes it clear that we have not so narrowly restricted our definition of position that substantial modifications to our existing conception of it become impossible. The second answer is that we should distinguish between, on the one hand, learning something new about an existing physical concept — in which case *something* must make it ‘the same’ concept both before and after we learned something new about it (and we are proposing that what is the same is its relationship to certain parts of the Galilean group) — and, on the other hand, discovering new physical concepts. We have no proposals to make here about how such discoveries occur, or how to understand them.⁴³

In the remainder of this subsection, I will discuss in further detail the claims made above, that insisting on the relevant invariances and covariance is sufficient to fix the position and momentum observables. In the next subsection, I will turn to the role of reference frames in the definition of observables.

3.3.2.1 Systems of Imprimitivity and the ‘Uniqueness’ of Quantum Observables Position and momentum, considered as POVMs, each give rise to a ‘system of imprimitivity’, a special case of a system of covariance.⁴⁴ Generically, a system of covariance is a set $(\mathcal{H}, E, S, \mathcal{G}, \alpha, \{U_g\})$ where \mathcal{H} is a Hilbert space, E is a POVM whose domain is S and whose range is positive operators on \mathcal{H} , \mathcal{G} is some group, α is the action of \mathcal{G} on S , and $\{U_g\}$ is a unitary representation of \mathcal{G} on

⁴³Tanona [2006] is particularly helpful on this point.

⁴⁴See Landsman, Ch. 5, this volume, for further discussion of imprimitivity and its application to issues in quantum theory.

\mathcal{H} . If E is a PVM then the set is a *system of imprimitivity*.⁴⁵ Systems of imprimitivity have important properties, in part summarized by Mackey's imprimitivity theorem. My discussion here will follow that of Mackey [1996], emphasizing the structure and assumptions of the argument, rather than the mathematical details, which can easily be found in many places.⁴⁶

In a system of imprimitivity, S is often taken to be quite generic — e.g., it might be a generic metric space with \mathcal{G} some locally compact and separable group of isometries (assumed to have a continuous and transitive action on S — see §7.6.7). However, we will immediately specialize, in order to make quicker and more evident contact with real physical concerns. With that goal in mind, it is natural to take S to be \mathbb{R}^3 and \mathcal{G} to be, for example, the semi-direct product (§7.6.2) of translations and rotations ($\mathcal{A} \times \mathcal{R}$). However, in preparation for an application of the imprimitivity theorem, it is more useful to take S to be the topological group of translations ($\mathcal{A} = \mathcal{G}/\mathcal{R}$), which is clearly isomorphic to \mathbb{R}^3 as a topological space (indeed, as a metric space, given a suitable, and obvious, metric on \mathcal{A}). The idea, then, is that elements of S represent ‘displacements’ from some fixed origin, and thereby represent a position (and so are possible values of a position observable). The subgroup \mathcal{R} describes rotations around this origin.

We now require that a PVM for position, E^Q , be covariant with respect to $\mathcal{G} = \mathcal{A} \times \mathcal{R}$. (The action of any $g \in \mathcal{G}$ on $S = \mathcal{A}$ is defined in the obvious way: for $g \in \mathcal{A}$ its action on $a \in \mathcal{A}$ is just ga ; $g \in \mathcal{R}$ acts as the identity on \mathcal{A} . The action thus defined is transitive.) Finally, given a representation $\{U_g\}_{g \in \mathcal{G}}$ of \mathcal{G} on some Hilbert space, \mathcal{H} , we have a system of imprimitivity.

In one form, the imprimitivity theorem is the following.

Theorem (Mackey): Let $\{U_g\}_{g \in \mathcal{G}}$ be a unitary representation of a separable, locally compact, topological group, \mathcal{G} , on a separable Hilbert space, \mathcal{H} , and let \mathcal{K} be any closed subgroup of \mathcal{G} . Let E be a PVM whose domain is \mathcal{G}/\mathcal{K} such that $(\mathcal{H}, \mathcal{G}/\mathcal{K}, \mathcal{G}, \alpha, \{U_g\}, E)$ is a system of imprimitivity (with α a transitive action of \mathcal{G} on \mathcal{G}/\mathcal{K}). Then for any representation, $\{V_k\}_{k \in \mathcal{K}}$ of \mathcal{K} on some Hilbert space, \mathcal{H}' , the representation of \mathcal{G} induced by $\{V_k\}$ on the Hilbert space $L^2(\mathcal{G}/\mathcal{K}) \otimes \mathcal{H}'$ (and this induced representation exists) is unitarily equivalent to $\{U_g\}$. Moreover, $E(\Delta)$ (or its appropriate unitary transform, if the induced representation is related to $\{U_g\}$ by a non-trivial unitary transformation) must be the tensor product of the multiplication operator χ_Δ on $L^2(\mathcal{G}/\mathcal{K})$ with the identity on \mathcal{H}' .

⁴⁵Systems of covariance can be ‘dilated’ to systems of imprimitivity via the Neumark dilation theorem. See, e.g., Cattaneo [1979]. It follows, in essence, that we lose no generality by considering systems of imprimitivity.

⁴⁶In addition to Mackey [1996], see Busch, Grabowski and Lahti [1995] and references therein for additional mathematical details and alternative routes to the same conclusion. Full details are available in Varadarajan [1985].

So in the case we are considering, we will let \mathcal{H} be $L^2(\mathbb{R}^3)$.⁴⁷ Then let $\{D_r\}_{r \in \mathcal{R}}$ be a representation of \mathcal{R} on some Hilbert space, \mathcal{H}' , and consider the associated induced representation of \mathcal{G} . This induced representation must be unitarily equivalent to *any* representation of \mathcal{G} , and moreover, $E^Q(\Delta)$ is just $\chi_\Delta \otimes \mathbb{I}'$.

Now consider the simplest case, where $\{D_r\}$ is the trivial identity representation (i.e., every element of \mathcal{R} is represented by the identity on a 1-dimensional Hilbert space). Mackey's theorem immediately yields the usual Schrödinger representation of the position operator.⁴⁸ More explicitly, for $\Delta \in \mathbb{R}^3$, $E(\Delta) = \chi_\Delta$, where the latter is the operator with action 'multiply an element of $L^2(\mathbb{R}^3)$ by the characteristic function of Δ '. (This operator is a projection.) Integrating over all of S with respect to this PVM yields the usual position operator:

$$(80) \quad Q = \int_{\mathbb{R}^3} \vec{r} dE^Q(\vec{r}),$$

which is a 'vector' of operators, Q_x, Q_y, Q_z , with the action, for $\vec{r} = (x, y, z)$, $Q_x \psi(\vec{r}) = x \psi(\vec{r})$ for any $\psi(\vec{r}) \in L^2(\mathbb{R}^3, d\vec{x})$, and similarly for Q_y and Q_z . Notice what has happened, here. We began with the requirement that the PVM representing position have the 'correct' covariances with respect to translations and rotations — i.e., that it be a part of the relevant system of imprimitivity — and we ended up, via Mackey's theorem, showing that up to unitary transformation, we *must* choose the usual ('Schrödinger') representation of position. Another, somewhat more operational, way of putting the point is this: assuming that observations (or predictions) involving the position observable respect the symmetries (translations and rotations) of the space of possible values of position, all representations of the position operator (PVM) are unitarily equivalent to the Schrödinger representation. In terms of establishing the empirical meaning of the theory, we might say that what it *means* to 'be the position' of a system is to transform in the right way under the action of the relevant pieces of the Galilean group. Mackey's theorem establishes the uniqueness (up to unitary transformation) of position, so defined.

Indeed, Mackey's theorem establishes more. It also establishes the relationship between the representation of (i.e., choice of a POVM for) position (and momentum) and the representation of the Galilean group. To see why (in outline), let \mathcal{G} be any group continuously parametrized by $a \in \mathbb{R}$ (§7.6.4). Now, as we have already discussed, if \mathcal{G} is supposed to be a *symmetry* group, then in general it should be represented, quantum-mechanically, in terms of transformations of the Hilbert space that 'make no difference, physically', and such transformations are often given in terms of unitary operators. (Recall the discussion from §1.5.1.2.) Moreover (recall §1.5.1.4), when the group is continuously parametrized by a (as it will be in the cases of interest for us), these unitary operators are given in terms

⁴⁷More precisely, we should use the topological group of translations in place of \mathbb{R}^3 , but we already noted that for our purposes these are the same spaces.

⁴⁸The result is a description of a spinless particle — see Mackey [1996]. Non-trivial representations of \mathcal{R} result in the description of particles with spin $\frac{\dim(\mathcal{H}')-1}{2}$.

of a self-adjoint operator, F , on \mathcal{H} such that $U_a = e^{-iF a}$. Hence, for example, the unitary representation of spatial translations and boosts must take this form. Mackey's theorem then implies that, in fact, translations and boosts are given by (78) up to unitary equivalence.

It is important to keep in mind that 'up to unitary equivalence' does *not* mean that one can apply *different* unitary transformations to the U_a and V_b (from eq. 78) and still satisfy all of the (invariance and covariance) conditions that have been placed on the POVMs for position and momentum. The point here is that the *relationship* between position and momentum is established by those conditions via Mackey's imprimitivity theorem. (On the other hand, one can always apply a *global* unitary transformation, but such a transformation is akin to 'translating the universe five feet to the right' in classical mechanics.)

Indeed, one can establish, from these results, the conclusion that position and momentum must obey the Weyl form of the commutation relations.⁴⁹ In particular,

$$(81) \quad U_a V_b = e^{iab} V_b U_a, \quad (a, b \in \mathbb{R}).$$

As we will discuss later (§4.1.2), this expression is a version of the canonical commutation relation (24) between P and Q . In other words, the assumption that position and momentum bear the right relation to the Galilean group leads directly to their incompatibility.

Hence the role that the position and momentum *operators* (POVMs) play in the Galilean group, the action of the Galilean group on them, and the (Weyl form of the) commutation relations between them, are fixed as soon as we insist on two things: (1) that position and momentum satisfy the invariance and covariance conditions given above, and (2) that position and momentum, as well as the Galilean group itself, be represented on a Hilbert space. There seems to be little room to deny (1) in a non-relativistic theory, and (2) may be understood as the requirement that our theory be genuinely quantum-mechanical.⁵⁰ Finally, one should bear in mind the *generality* of Mackey's theorem. I have discussed it largely in the context of position and momentum, but analogous accounts hold for any observables (POVMs) that form a system of imprimitivity (or covariance) with regard to some group of symmetries. Some examples (among them, angular momentum and spin) can be found throughout Busch *et al.*, [1995]. We leave the

⁴⁹Consult Varadarajan [1985, ch. V] or Mackey [1949; 1978] for details.

⁵⁰We must be a little careful, here. (Thanks to Jos Uffink for raising this question.) It is well known that classical mechanics can be represented as a theory on a complex Hilbert space. See Bracken [2003] and references therein, the originator of the idea being, apparently, Groenewold [1946]. However, the classical observables on Hilbert space form (unsurprisingly) a commutative algebra, in virtue of the definition of a 'non-standard' product of linear operators (the 'odot' product in Bracken [2003]), one that does not make an appearance in quantum theory. (The dynamics, for example, are defined in terms of a Lie bracket defined in terms of this 'odot' product rather than in terms of the usual product of operators given by the composition of their action on the space.) Hence the more careful statement of the claim in the text is that we require position and momentum to be operators in the algebra of operators on a Hilbert space *as standardly understood*, that is, under the usual product of operators (composition).

reader to investigate, and turn now to consider the role of reference frames in the definition of physical quantities in quantum mechanics.

3.4 Reference Frames

3.4.1 Identification and Role of Reference Frames in Quantum Theory

It has been claimed (not unreasonably) since the early days of quantum theory that there is no room in the theory for the notion of a reference frame. The apparent difficulty is this: a reference frame, *by definition*, has a well-defined location and state of motion, because locations and states of motion are defined relative to it. But then it is hard to see how there could be any such thing as a quantum reference frame, because, as I have mentioned already (§1.2.5), and shall discuss in some detail below (§4), standard quantum theory cannot describe anything as having a well-defined location and state of motion (momentum). Bohr (on one reading⁵¹) concluded that it is up to *us* to *stipulate* some object (normally, a measuring apparatus) as defining a reference frame, and that this stipulation requires us to treat the object classically, because the stipulation requires the object to be well-defined in position and momentum. (Of course, we can also step back and describe that object quantum-mechanically again, having stipulated some *other* object to serve the role of a reference frame.)

However, there are good and bad stipulations — as Bohr himself emphasized, not every object is reasonably taken to define a reference frame for a given purpose. One important reason is that there is a *prima facie* distinction between inertial frames and reference frames. Inertial frames are the frames in which the laws of motion (whether classical or quantum) are valid.⁵² It is a matter for empirical investigation to seek out inertial frames. Reference frames, on the other hand, are the frames used to define physical quantities, such as position, momentum, angular momentum, spin, and so on. For the instantaneous definition of these quantities, *any* stipulation of a reference frame is fine.

However, the choice of a reference frame for the description of systems *over time* is far from arbitrary. For example, a rotating coordinate system used as a reference frame will introduce fictitious Coriolis forces. Of course, by ‘rotating’ one means ‘rotating relative to an inertial frame’, and herein lies the crucial point: a reference frame that is not inertial will always introduce fictitious forces (i.e., apparent violations of the laws of motion). Hence, in the end, although it is certainly permissible to describe our physical systems with respect to whichever reference

⁵¹See Bohr [1935]. For some (albeit flawed) interpretive remarks on this paper, see Dickson [2002a; 2002b]. Recent detailed interpretations of Bohr’s insistence on the necessity of classical concepts can be found, for example, in Tanona [2002] and Howard [2003]. Of course, many others have also written on this topic. The references in those works will get the interested reader started.

⁵²There is a long history behind this understanding of what it means to be inertial. See DiSalle [1990; 2002] and Barbour [1989]. It culminates in the idea that Newton’s law of inertia should be understood as the claim that there *is* a reference frame, an ‘inertial frame’, in which the other two laws are true. This idea can be extended to quantum theory.

frame is most convenient, it is also necessary that we know *how* to describe the system in terms of a reference frame that is inertial, by which we mean, here, one in which the laws of motion — classical or quantum, as the case may be — are true.

Nothing in quantum theory rules out the possibility of an inertial frame in this sense. Indeed, just as classical physics does, quantum physics contains an assumption (usually left implicit) that there is some frame (some system of coordinates) in which the laws are valid. (It does not follow that such a system of coordinates can be used as a reference frame in the *classical* sense, i.e., to define, simultaneously, definite positions and momenta.)

In quantum mechanics, as in classical mechanics, one goes about finding an inertial frame by searching for coordinates in which the dynamical laws are true. In classical mechanics, this search often extends to frames defined in terms of celestial bodies (Ma *et al.*, [1998]). In quantum theory, one does not go to such lengths. Typically measuring devices, or the labs that contain them, are sufficient to serve the purpose of defining a reference frame.

3.4.2 Relational and Absolute Coordinates

If it is correct that quantum observables are typically defined relative to a reference frame (whether inertial or otherwise), then typically, in quantum theory, the observationally significant variables will be relational in character.

This point does not apply only to the case of measurements involving position (or $L^2(\mathbb{R}^{3N})$ spaces) explicitly. Consider spin. What, for example, does it mean *observationally* to say that a system is in the state $|\psi\rangle = |z_+\rangle$? If we are not told *which* direction in space counts as z then the claim that a system is in that state is observationally empty.

However, we saw above that when we wish to apply the quantum laws of motion to a system, we must do so in coordinates given by some frame in which the laws are valid. We also mentioned that in quantum theory, this frame is often determined by some macroscopic piece of apparatus. What are the conditions that must be met by this macroscopic piece of apparatus in order for it to serve the purpose of defining coordinates in which the laws are valid? And how are we ever in a position to verify that those conditions are met?

In the end, the answer to the last question is the same as in classical mechanics — we never have access to any ‘absolute’ frame, one that is known to be inertial, from which we can check the inertiality of other frames (and hence their suitability to define coordinates in which the laws are valid). The best we can do is to determine, empirically, as best we can, that the laws are valid in some particular frame, F , and then justify the use of other frames by reference to F , and in particular by noting that these other legitimate frames are related to F by an appropriate symmetry transformation.

In other words, at least until we have a truly relational theory on hand, the relational coordinates are ultimately defined in terms of the absolute coordinates,

and they are so defined in terms of certain symmetry transformations, as we shall now describe.⁵³

Let us begin by recalling that the transformations from one inertial frame to another are (in the present non-relativistic context) given by the Galilean transformations. (Moreover, we have seen that we have little choice about the mathematical form that these transformations take in quantum theory.) We can use that fact to derive an expression for a transformation from the absolute coordinates of some given inertial frame to coordinates measured relative to some stipulated reference frame.

A helpful way to consider the situation is as follows. Imagine an observer, A , inside a lab and suppose that A measures physical quantities relative to the lab. Now imagine an ‘external’ observer, B , who has been *given* the information (or has assumed it) that some frame, F , is inertial. B also uses F as a reference frame: as far as B is concerned, the lab and its contents are all described in the coordinates given by F . But suppose B wishes to describe A ’s measurements *as relational and in the frame given by the lab*. How does B transform from the coordinates given by F to the (relational) coordinates given by the lab (which could also be moving relative to F)?

The answer follows more or less immediately from the form of the Galilean transformations. (Of course, we are not simply applying a Galilean transformation in this case. We are also transforming to relational coordinates.) As Aharonov and Kaufherr [1984] point out, the correct transformations are

$$(82) \quad U_{AK} = e^{-i \sum_{n>0} P_n^B Q_0^B}$$

where P_n^B represents the momentum observable used by B to describe the momentum of system n , and Q_n^B similarly represents B ’s position observable. System 0 is the lab itself. (Note that their result implicitly assumes that the reference body — the ‘lab’ — moves inertially in F . Notice also that in this situation, the lab drops out of view. A has no coordinates to describe the lab (system 0) because A ’s coordinates are all defined *relative to* the lab.)

As a quick check, note that

$$(83) \quad U_{AK} Q_n^A U_{AK}^{-1} = Q_n^B - Q_0^B \quad (n > 0),$$

so that, as expected, what A describes but Q_n^A , B ‘knows to be’ $Q_n^B - Q_0^B$. That is, B can describe, in purely quantum-theoretic terms, the fact that A ’s measurements of position are made relative to the lab.

The main lesson of this discussion, for us, is this: in general, we are in the position of observer A , not that of B . We are not *given* an absolute frame. Instead, our quantities are measured (and therefore, operationally at least, defined) relationally. Nonetheless, the symmetries that take us from one frame to another are ultimately involved in the definition of those quantities, because in order for our physical theories legitimately to apply to them, we must consider them to be

⁵³The discussion here is motivated primarily by the work of Aharonov and Kaufherr [1984].

‘really’ defined in terms of the coordinates of an ‘absolute’ inertial frame, and such a definition involves the symmetry transformations that take us from one (inertial) frame to another.

An interesting side-point here is the following. Observer B (whose point of view is, in principle, the truly legitimate one, here), has a particularly interesting way of noticing the incompatibility of position and momentum. Suppose A is going to measure the position of a particle, and writes down an interaction Hamiltonian along the lines of the model of impulsive measurement (§3.1.2). Of course, A just writes down something like $H_I = g(t)\Pi Q$, where Q is A ’s position operator for the measured system. When B transforms this Hamiltonian to the *correct*, i.e., *relational* coordinates, and solves the equation, B finds that as a result of the measurement, the *lab itself* experiences a shift in momentum, which renders the lab unsuitable (for A) for defining momentum (because A has no way to measure the shift in the momentum of the lab — the lab is the reference frame relative to which any such measurement could be made, for A).⁵⁴

3.5 A Group-Theoretic Characterization of Empirical Content

3.5.1 Reframing the Issue

How does the discussion of §3.3 and §3.4 help to establish an account of the empirical content of the quantum formalism? Of all of the questions raised in §3.2, let us focus on the following two: (1) What empirical meaning is attached to the elements in the domain of a POVM? (2) Given an answer to (1), how is the POVM itself chosen appropriately?

The discussion above suggests the following general approach. (a) Observables are frame-dependent quantities that are *defined* in terms of their behavior with regard to some group of symmetries. (b) Having pointed out that legitimate reference frames ought to be related in the right way to inertial frames — more generally, they ought to be suitable for defining the quantities that we indeed wish to define — we ought then seek to attach some empirical, observational, significance to such things. (c) Similarly, having pointed out the role that symmetry transformations play in the definition of physical quantities, we ought then to seek to attach some empirical, observational, significance to such things. A more detailed version of (a) would constitute an answer to question (2) above. More detailed versions of (b) and (c) would constitute an answer to question (1) above.

In fact, we have already said something about (a), and considerably more has been said in the literature about the symmetries that are obeyed by (and thus that are definitive of) quantum-mechanical observables. In any case, the basic point is that the (reference-frame-dependent) observables may be uniquely characterized

⁵⁴See Dickson [2004b] for details. It is not sufficiently emphasized there that these observations are at best a first step towards understanding the incompatibility of position and momentum. Note, moreover, that nothing in B ’s calculations implies the uncertainty relations, understood as placing a lower bound on the precision with which position and momentum can be simultaneously measured, or known, or defined.

(up to trivial transformations, such as scaling of length, and global unitary transformations) by the invariances and covariances that they obey with respect to some appropriate symmetry groups. In this sense, the very meaning of ‘position’, ‘momentum’, and so on, is partially given by these invariances and covariances. I conclude this section with some preliminary thoughts about (b) and (c), followed by the consideration of an objection.

3.5.2 *The Empirical Content of Frames and Transformations*

Concerning (b), the basic proposal here is that reference frames describe the world as witnessed by some observer, and ultimately, some human being. This proposal reflects a point of view according to which theories are, ultimately, human constructions. This point of view does not (necessarily) include the idea that there is any (or much) arbitrariness in physical theory — the world might still dictate how observational creatures such as ourselves are bound to construct theories, if we are to be successful. It does entail that the ‘observables’ of a theory are intimately connected with the observational capacities of human beings, and with the properties of those capacities.

However, those who would subscribe to such a view must be careful, for at least two related reasons. First, reference frames are typically idealized in various ways that might not apply to actual human observation. They are, for example, typically taken to be entirely rigid (spatially). Second, if a given reference frame is to be used over any stretch of time, then either it must be inertial, or one must know how to relate it to an inertial frame, in order legitimately (and successfully) to apply the law of motion (whether that be Newton’s second law, Schrödinger’s equation, or something else). But as noted earlier, in actual practice it is very difficult to determine whether a given frame is inertial. Nonetheless, the proposal being floated here is that ultimately reference frames should be understood as ‘legitimate (and idealized) human points of view’, where the notion of legitimacy is to be spelled out in terms of a known connection with an inertial frame.

Concerning (c), the empirical content of the symmetry transformations themselves, I again offer a kind of anthropocentric view. The suggestion, coming from various 19th century philosophers of geometry (for example Helmholtz, and in a different way, Poincaré), is that such transformations are connected to experience via the physiological-kinesthetic experience of undergoing the change from observing from one reference frame, to observing from another. (Consider, for example, the physiological-kinesthetic experience associated with a rotation, or spatial translation.) However, the connection of specific groups of transformations with empirical content of various sorts is far beyond the scope of this essay. We merely note that an essential ingredient in the program outlined here is establishing (or understanding) such connections, for the groups that are at the heart of quantum theory.⁵⁵

⁵⁵One could adopt an approach completely opposite to the one suggested here, one that takes a theory as somehow ‘already’ empirically meaningful *prior* to the empirical account of the relevant

3.5.3 ‘Absolute’ Quantities

What about quantities that are *not* reference-frame dependent? Surely one of the lessons of relativity theory is that while many measured quantities are frame-dependent, some physical quantities — perhaps even the most important ones — are ‘absolute’, i.e., frame-independent. Think, for example, of the spacetime interval, defined by $\tau^2 := t^2 - (x^2 + x^2 + x^2)$. It is frame-independent, in the sense that *whatever* coordinates one uses to calculate τ , the result is always the same. Indeed, one understanding of *non*-absolute quantities is that they are nothing more than the absolute quantities seen ‘from a particular perspective’, so that the absolute quantities are somehow fundamental, the frame-dependent ones derived.

But how does one *measure*, or *observe* the value of an absolute quantity? The claim made here is that we cannot help but do so from within some reference frame. While the *result* is not dependent on the reference frame, the measurement still occurs in one. (Consider, for example, how one might measure τ .) If this claim is correct, there there are two responses to the point that some quantities are absolute.

The less radical response is to allow that the absolute quantities might even be more fundamental in some sense. However, we are here concerned with the *observational* content of quantum theory — how does the mathematical formalism connect with experimental observation? If, as claimed above, observation always takes place within a frame, then it is fair enough for an account to consider the empirical content of just the frame-dependent quantities.

The more radical response is to assert the reverse of the point of view mentioned above: the fundamental quantities are the frame-dependent ones, and the ‘absolute’ quantities are derived (calculated) from them. This view is in fact closely connected with the view, expressed above, that (our) physical theories are fundamentally about the world *as observed by us*. On this view, the role of the absolute quantities is to underwrite the possibility of communication amongst inhabitants of different frames about the values of their frame-dependent quantities; it makes possible some sort of agreement; and that is all.

4 UNCERTAINTY

This section is devoted to an examination of uncertainty in quantum theory. We will begin (§4.1) with the formal origins of the uncertainty relations, namely, the canonical commutation relations. We will then consider the uncertainty relations (§4.2), which, minimally, express the fact that the dispersions of two non-commuting observables (a notion that we will define) cannot be made simultaneously arbitrarily small. We will then (§4.3) consider two fundamentally different ways to understand, or interpret, the uncertainty relations, and finally (§4.4), we will consider in some detail the famous argument by Einstein, Podolsky, and Rosen

groups of transformations, and then define the empirical content of those transformations in terms of their consequences within the theory. I shall not follow that idea through here.

[1935], intended to cast some form of doubt on the fundamentality of quantum uncertainty, and ultimately on the ‘completeness’ of quantum theory as a description of physical reality.

4.1 Canonical Commutation Relations

4.1.1 Representations of the Commutation Relations

Earlier (§1.2.5) we note that the quantum-mechanical position (Q) and momentum (P) operators obey the CCRs (24). It is perhaps more accurate to say that in quantum theory, Q and P are *chosen* so that (24) is obeyed. Choosing operators, Q and P , on some Hilbert space, \mathcal{H} , such that (24) is satisfied is choosing a *representation* of the CCRs. It is, arguably, the satisfaction of these commutation relations by (the operators that stand for) certain pairs of observables — centrally, position and momentum — that makes a theory truly ‘quantum’.

It turns out that for any representation, the operators cannot both be bounded, and therefore \mathcal{H} must be infinite-dimensional (§7.2.2). Heisenberg constructed a representation on the space ℓ^2 (§1.2.1) in terms of infinite-dimensional square matrices:

$$(84) \quad Q = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 & 0 & \dots \\ 1 & 0 & \sqrt{2} & 0 & \dots \\ 0 & \sqrt{2} & 0 & \sqrt{3} & \dots \\ 0 & 0 & \sqrt{3} & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \quad P = \frac{-i}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 & 0 & \dots \\ -1 & 0 & \sqrt{2} & 0 & \dots \\ 0 & -\sqrt{2} & 0 & \sqrt{3} & \dots \\ 0 & 0 & -\sqrt{3} & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Schrödinger constructed a representation in terms of operators on the space $L^2(\mathbb{R})$ in which Q is the multiplication operator (i.e., $Qf(x) = xf(x)$ for any $f \in L^2(\mathbb{R})$) and $P = -i \frac{d}{dx}$.

These two representations are in fact unitarily equivalent. (I.e., there exists an isomorphism from ℓ^2 to $L^2(\mathbb{R})$, under which the Heisenberg operators go to the Schrödinger operators.) Any representation that is isomorphic to these is called *regular*. Non-regular representations exist.⁵⁶

Note that because (at least one of) the operators in a representation of the CCRs must be unbounded, we must be careful about keeping track of their domains of definition (§7.2.2). The CCRs are thus defined only on some (dense) subset of the space.

4.1.2 The Weyl Relations

The fact that P and Q must be unbounded is occasionally bothersome. For example, we just noted that this fact requires one to keep track of their domains of

⁵⁶Here is an easy example. Consider the space of square-integrable functions on the open interval $(0, 1)$. Let Q and P again be the multiplication and differentiation operators (as given above). In this case, Q is in fact bounded (but P is not), and so this representation cannot be unitarily equivalent to the Schrödinger representation, because, in that case, both of the operators are unbounded.

definition. An alternative approach, due to Weyl, avoids the problem. We begin by considering a pair of strongly continuous one-parameter unitary groups of operators U_a and V_b , $a, b \in \mathbb{R}$ (§7.6.4). We will call them a *Weyl pair* if they satisfy the relation (81) from §3.3.2.1. By Stone's Theorem (§1.5.1.4), U_a and V_b can be written as

$$(85) \quad U_a = e^{-iaQ}, \quad V_b = e^{-ibP}$$

where Q and P are unbounded selfadjoint operators, defined on a common (dense) domain. (Recall eq. 78.) Writing these exponentials formally, in terms of a power series expansion,

$$(86) \quad e^{-iaQ} = \sum_{n=0}^{\infty} \frac{(-iaQ)^n}{n!},$$

(similarly for e^{-ibP}) and substituting into (85) we retrieve the CCRs. (If both sides of (85) are defined on a common dense subspace, then this procedure is rigorously valid. Otherwise, it is suggestive symbol-manipulation.) Note, finally, that the operators e^{iaQ} and e^{ibP} are bounded, so that, for example, issues about the domain of definition do not arise.

4.1.3 Von Neumann's Uniqueness Theorem

There is another nice feature of the Weyl relations: all representations of them are regular. In other words, every Weyl pair is unitarily equivalent to the Weyl pair generated by the Schrödinger position and momentum operators. This result, due to von Neumann [1931], implies that the generators of any Weyl pair must have spectra that are the entire real line.

A further question concerns when a representation of the CCRs is regular — i.e., how can one tell, apart from checking for unitary equivalence directly (which can be hard) that a given representation is regular? Answers to this question (beyond the observation just made about spectra) are known (Rellich [1946], Dixmier [1958], but beyond the scope of this essay). See Summers [2001] for discussion and further references.

4.2 The Uncertainty Relations

One of the reasons that the CCRs are so important is that they give rise, directly, to the uncertainty relations. Recall that if two bounded operators do not commute, then there are eigenvectors of one that are not eigenvectors of the other. (Similar remarks hold for unbounded operators, but there we must take account of the fact that they might not have any eigenvectors, and speak instead in terms of the non-commutativity of their spectral projections.) It follows that there are states that assign trivial probabilities to the possible values of one observable (i.e., probability 1 for one eigenvalue, and 0 for the others), and non-trivial (not 0 or 1) to at least two possible values of the other. Hence non-commutativity already

implies a type of ‘uncertainty relation’: certainty about the value of one observable can imply uncertainty about the value of another. Below, we shall make this idea more precise, and consider its interpretation.

4.2.1 *The Optical Derivation*

In [1927], and in an improved version in [1930], Heisenberg made the following argument, intended to make some sense of, perhaps even to derive, the uncertainty relations for position and momentum. Suppose we wished to measure the location of a small particle (e.g., an electron), by means of an optical microscope. The resolving power of the microscope with an aperture angle θ is approximately $\lambda/\sin\theta$, where λ is the wavelength of the light. This resolving power determines our uncertainty about the position of the particle after the measurement. On the other hand, in order for us to detect the particle, at least one photon would have to strike it. This photon has momentum⁵⁷ h/λ and the angle of impact is uncertain to within the angle θ ; hence the amount of momentum transferred to the particle is uncertain to within roughly $(h\sin\theta)/\lambda$, and the *product* of the uncertainty in the position and momentum of the measured particle, after the measurement, is roughly h .

In other words, there is a lower limit on the product of the uncertainty of position and momentum. Notice that this lower bound applies only *after* the measurement. Indeed, we could measure the momentum of the particle with arbitrary precision prior to measuring its position, also with arbitrary precision. We would then have determined its position *and* momentum just *prior* to the moment of impact (of the photon on the particle) with arbitrary precision (though we would still be uncertain about its momentum *after* the measurement).

There are other derivations of the uncertainty relations, derivations that rely more explicitly on the formalism of quantum theory. We now consider two of them.

4.2.2 *The Wavefunctional Derivation*

Position and momentum are related by a Fourier transformation. Indeed, it is often easier, when dealing with momentum in the Schrödinger representation, to apply a Fourier transform (which is a unitary transformation on the Hilbert space), so that the momentum operator becomes a multiplication operator (and the position operator then becomes differentiation). But consider what happens to the wavefunctions. A wavefunction that is very well peaked corresponds to a state in which position is well-defined. That is, most of the probability is concentrated in a relatively small region of space. But the Fourier transform of such a function is very flat, so that the probability is uniformly distributed across all of the real line, which, after the transformation, corresponds to possible momenta of the particle.

⁵⁷ The relation $p = h/\lambda$ as applied to photons was introduced as part of Einstein’s [1905] explanation of the photoelectric effect, and generalized to material particles by de Broglie [1924].

This general idea can be made mathematically more precise. Consider a Gaussian ‘wavepacket’, a wavefunction from $L^2(\mathbb{R})$, which, as a function of x , has appreciable magnitude only in some region of size $2a$:

$$(87) \quad \psi(x) = e^{-x^2/2a^2}.$$

The Fourier transform of this wavefunction (i.e., transforming to a ‘momentum representation’) is:

$$(88) \quad \bar{\psi}(k) = e^{-a^2k^2/2}$$

(where k is the ‘wave number’; momentum is given by $p = \hbar k$). This Gaussian has width $2/a$. Hence a narrowly peaked wavefunction in position (i.e., a is small) is widely spread in momentum (i.e., $1/a$ is large). Specifically, setting $\Delta x \approx a$ and $\Delta p = \hbar \Delta k \approx 1/a$, we have that $\Delta x \Delta p \approx \hbar$. (This expression is not quite the standard uncertainty relation, but then the derivation here is not meant to be exact.)

4.2.3 The Algebraic Derivation

There are in fact many roads from the quantum formalism to the uncertainty relations. Here, we consider just one other commonly found derivation, in part because it will shed some additional light on meaning of ‘uncertainty’, and in part because unlike the previous two derivations, this one is rigorous, and results in the exact form of the uncertainty relations.

Given an observable, F , define $\Delta F := F - \langle F \rangle$. (The right-hand side is the expectation value of F , where we have left the state unspecified.) The expectation of $(\Delta F)^2$ is the *dispersion* of F . Indeed,

$$(89) \quad \langle (\Delta F)^2 \rangle = \langle F^2 - 2F\langle F \rangle + \langle F \rangle^2 \rangle = \langle F^2 \rangle - \langle F \rangle^2,$$

which is a standard statistical notion of ‘dispersion’ (often called ‘mean-square deviation’ by physicists, and ‘variance’ by statisticians; its square root is the standard deviation). Now, let F and G be observables (self-adjoint operators). Then the Schwarz Inequality (§7.1.3) implies that

$$(90) \quad \langle (\Delta F)^2 \rangle \langle (\Delta G)^2 \rangle \geq |\langle \Delta F \Delta G \rangle|^2.$$

Straightforward algebraic manipulations (see, e.g., Sakurai [1985, p. 36]) transform (90) into the standard *uncertainty relation*:

$$(91) \quad \langle \Delta F \rangle \langle \Delta G \rangle \geq \frac{1}{2} |\langle [F, G] \rangle|.$$

for any observables F and G . Notice, for example, that

$$(92) \quad \langle \Delta P \rangle \langle \Delta Q \rangle \geq \frac{1}{2}.$$

(or, if we are not setting $\hbar = 1$, then the right-hand side is $\hbar/2$ — see note 11).

Above, we said that this derivation would shed some light on the meaning of uncertainty. In particular, we can now see that, rigorously, ‘uncertainty’ refers to the dispersion (standard deviation) of an observable in a given state, which is normally understood as a measure of the ‘spread’ in its values over an ensemble of systems all in some given state.

4.2.4 *Limitations and Generalizations*

The derivation of (91) makes it clear that *any* two non-commuting operators will give rise to some uncertainty relation. Hence (91) is quite general. However, there are also ‘uncertainty relations’ between quantities that are not represented by operators. The most well-known is the time-energy uncertainty relation, whose interpretation is notoriously problematic precisely because time is not an observable in quantum theory. (There is no self-adjoint operator that represents time.) Another example is phase and photon number. (Again, there is no ‘phase’ operator.) Various proposals exist for how to understand these other uncertainty relations, but here we merely note the point that they must apparently be understood in some sense other than that given by (91).

Moreover, (91) faces other problems. First, recall that (91) is state-dependent. Indeed, if we choose a state that is an eigenstate of F (or G), then even if F and G do not commute, both sides of (91) are zero, which certainly seems to violate the spirit of ‘uncertainty’. (The ‘solution’, perhaps, is to notice that if F and G do not commute, then in general the dispersion of G for a system in an eigenstate of F will be non-zero.) Second, keep in mind that ‘dispersion’ itself can be misleading. Even when most of the probability is concentrated on a narrow range of possible values of F , a small amount of probability that is located very far from F ’s mean can cause its dispersion to become large. There are proposals to handle these shortcomings. (See especially Uffink [1994].)

4.2.5 *‘Wave-Particle Duality’*

A quantum-theoretic experiment that is commonly associated with the uncertainty principle is the double-slit experiment (which had been done, in some form, from well before the advent of quantum theory). The set-up is as follows: a source of particles (or monochromatic light, i.e., photons) is placed in front of an opaque barrier with two parallel slits. Behind the barrier is a screen (such as a photographic plate). A particle is fired at the slits in the barrier, and the screen records the location of each particle as it strikes the screen. (See figure 3.)

The main point is the following. We shine a beam of particles from the source onto the barrier. If both slits are left open and no determination is made about which slit the particle traverses, then an interference pattern develops on the screen, which is just what one would expect if a wave somehow passes through the slits. (See figure 4a.) If, on the other hand, we determine which slit the particles traverse, then no interference pattern shows up on the screen — instead,




Figure 3. Double-Slit Experiment.

a pattern that is characteristic of particles (one ‘blob’ behind each slit) shows up. (See figure 4b.) What is more, one can do the experiment one particle at a time, and in this case, one sees ‘dots’ on the screen, and yet, if no determination is made about which slit the particles traverse, eventually the dots exhibit an interference pattern. (See figure 4c.)




Figure 4. Results of the Double-Slit Experiment: (a) with the particles passing freely through the slits; (b) while determining which slit each particle traverses; (c) with the particles passing freely through the slits, sent one at a time (simulated, after 800 particles are detected at the screen).

This experiment illustrates ‘wave-particle duality’: when we measure a wave-like property of particles (interference), we get wave-like behavior (interference pattern), while when we measure a particle-like property of particles (which slit a particle traverses), we get particle-like behavior (no interference pattern).

Indirectly, it also illustrates the uncertainty relations. For consider what it takes to determine with reasonable accuracy which slit a particle traverses — in that case, we must measure the particle’s position to an accuracy much better than $d/2$, where d is the distance between the slits. If the interference pattern is to be maintained despite this measurement, then the particle’s momentum cannot

be disturbed so much that, with appreciable probability, it gets deflected from a region of constructive interference (where, from the wave-theoretic point of view, the waves passing through each slit interfere constructively, i.e., a region where many dots show up in figure 4c) to an adjacent (or indeed any) region of destructive interference (i.e., regions where few or no dots show up in figure 4c). A rough trigonometric analysis shows that in fact the product of the uncertainty in our position measurement and the required low uncertainty in momentum must violate the uncertainty relation between position and momentum. In other words, the uncertainty relations appear to require that measuring the position of the particle at the slits well enough to determine (with good accuracy) which slit the particle traverses will tend to wash out the interference pattern (the more so, the more accurate the measurement of position).

4.3 Interpretation of the Uncertainty Relations

4.3.1 Observational-Epistemological

Interpretations of the uncertainty relations are helpfully divided into two kinds: those that understand uncertainty purely in terms of observationally obtainable facts about the values of observables, and those that attribute observational uncertainty to a more fundamental, ‘ontological’, uncertainty, or ‘indeterminacy’. I will consider each of these two kinds of interpretation in turn.

4.3.1.1 Uncertainty as Uncertainty The term ‘uncertainty’, and the understanding of uncertainty as standard deviation (dispersion), strongly suggests an epistemic, even operationalist, understanding of the uncertainty relations. In particular, one is quite tempted to understand the uncertainty relations in terms of post-observational uncertainty about the actual values of observables.

There is a question about whether this uncertainty concerns single systems, or ensembles. In the latter case — i.e., if, as suggested by the usual statistical understanding of standard deviation, we understand uncertainty to reflect dispersion of values (for some observable) in an ensemble — the ‘interpretation’ of the uncertainty relations does not go beyond a straightforward statistical understanding of standard deviation. Of course, in this case then there is still a sense in which uncertainty applies to single systems, namely, when they drawn at random from such an ensemble. But one might, in addition, suppose that the notion of uncertainty applies to single systems independently of the consideration of any ensemble. Indeed, one might argue that only in this case would we really *understand* why uncertainty holds at the statistical level. In any case, let us consider that possibility.

The idea, then, is that measurement (more generally, observation) in general reduces our uncertainty about the value of some observable, F , on the (single!) measured system. The uncertainty principle would then be read as asserting that a reduction in uncertainty about the value of F can imply an increase in uncertainty about the values of observables that are incompatible with F .

The meaning of such an assertion is clear enough, but understanding *why* incompatible observables have this feature is another matter. Suppose that F and G are incompatible. Suppose that we measure F . We then measure G . Why should it be the case that this second measurement *ruins* our previous knowledge of the value of F ? Indeed, supposing that the second measurement is isolated from the first, and that the value of F was not in any way disturbed in the meantime, *how* could the second measurement ruin our previous knowledge of the value of F ?

4.3.1.2 Einstein’s Early Thought Experiment Early critics of quantum theory — notably, Einstein — asked something like this question, and indeed proposed (thought) experiments that were apparently intended to show that in fact it *is* possible to ‘beat’ the uncertainty principle. A famous such experiment, due to Einstein, involves a standard two-slit apparatus (see figure 3) that is mounted on springs. The basic idea is to use the springs to determine which slit the particle traversed, without in fact disturbing the particle itself, by measuring the exchange of momentum between the particle and the barrier. (If the source is located on the plane exactly between the slits, then if the particle goes through the right-hand slit (and the slits are narrow enough), the barrier will in general experience a kick to the right, and so on.)

4.3.1.3 Reply to Einstein Einstein’s challenge to the uncertainty principle, here, is empirical: he is claiming that, contrary to what quantum theory allows, it *is* possible to prepare a system in a state of precise position and momentum (or at any rate, more precise than the uncertainty relations allow). The response must, therefore, be empirical. If indeed it is possible to prepare a system in a precise state of position and momentum, then we ought to be able to use our knowledge of either to make a verifiable prediction about the system. In particular, if the measurement of the position of the particle right after it passes through the slits does not, in fact, disturb the momentum of the particle, then we should still see the same interference pattern on the screen that we see in the standard two-slit experiment. If, on the other hand, the interference pattern ‘washes out’ (see §4.2.5), and approaches the ‘two blobs’ (see figure 4b) as the measurement of position becomes more and more precise, then the epistemic version of the uncertainty principle stands.

As far as current experiment can discern, it seems that a measurement of which slit the particle traverses does indeed wash out the interference pattern. Indeed, a remarkable experiment appears to show that not even this much is quite right; rather, what seems to matter is whether a *record* of the result of the measurement is kept, where by a ‘record’, here, we mean an encoding of the result in a measurable physical state of the universe. The experiment⁵⁸ is, in essence, the two-slit experiment, with a detector placed behind one of the slits. The detector is, moreover, ‘eraseable’, in the following sense. Once a particle passes through the detector, the particle leaves a trace in the state of the detector. We may then

⁵⁸See, for example, Scully and Walther [1989] and Walborn et al. [2002].

choose to magnify this ‘trace’ in order to turn it into a discernible signal indicating the presence of the particle, or we may completely erase it, so that the state of the detector no longer contains any retrievable information about whether the particle was once in the detector.

With this erasable detector in place, now imagine performing the following experiment. Fire the particles at the barrier one at a time. After the particle passes through the barrier, either erase the detector, or not. On the runs where we erase the detector, the particles build up an interference pattern (as in figure 4c). On the runs where we do not erase the detector, the particles do not build up an interference pattern, but instead ‘behave as particles’ (as in figure 4b). The experiment is relatively recent, and requires further scrutiny, but it strongly suggests, as do other quantum-mechanical experiments, that the epistemic version of the uncertainty principle is a fact of nature.

4.3.2 *Ontological*

With the epistemic version of the uncertainty principle apparently intact, we are left wondering *why* it is true. Different interpretations will propose different answers to this question. Some assert a ‘disturbance theory’, according to which the measurement of one quantity physically disturbs others (corresponding to incompatible observables) in an uncontrollable and unpredictable way. Others assert that reality matches our knowledge, here: we cannot know the value of G (when the value of incompatible F is known) because, in fact, G has no value.

There are many versions of this idea. Here we consider two. The first, often claimed to be part of the ‘Copenhagen’ interpretation of quantum theory, rests on some version of a verificationist, or operationalist, theory of the meanings of physical quantities, according to which a physical quantity has a value if and only if it has been measured (i.e., verified — the appropriate physical operations have been carried out, with the appropriate outcome). One must then argue that it is physically impossible to perform measurements of incompatible observables simultaneously on the same system.

This argument suggests a second ontological version of the uncertainty principle, according to which a physical quantity is well-defined just in case the conditions required for its well-definedness obtain. Of course, stated thus, this view sounds almost tautologous. It’s real content comes from arguing that there are indeed non-trivial physical conditions that must be in place in order for certain physical quantities to be well-defined, and that the conditions required for the well-definedness of a given quantity cannot be in place simultaneously with the conditions required for the well-definedness of any incompatible quantity.

This sort of argument has bite, if it does, because of the role of reference frames in the definition of physical quantities (§3.4). Hence, for example, ‘momentum’ must mean ‘momentum relative to X’, where ‘X’ is some physical system that defines a reference frame. But if ‘X’ (more precisely, the reference frame that it defines) is a non-inertial system, then it is not suitable for defining momentum (at

least not over any stretch of time), unless we know its relation to some inertial frame F (but then we are *really* defining momentum relative to F) — recall §3.4.1. Similarly, a non-inertial physical system is inappropriate for defining position (with the same caveat as before). These brief points recall the more extensive discussion above, and we shall have to leave the matter at that.

4.4 The Einstein-Podolsky-Rosen Argument

Thus far, everything that has been said about the uncertainty relations is consistent with the claim that the loss of certainty about the value of one observable, F , upon measurement of incompatible G , is a result of an unknown and uncontrollable physical disturbance of the value of F as a consequence of the measurement of G . Indeed, Heisenberg's 'optical' derivation of the uncertainty relations uses precisely this idea. But an argument due to Einstein, Podolsky, and Rosen (EPR)⁵⁹ purports to show that in fact the uncertainty relations cannot be understood in this way.

This argument concedes that the epistemic version of the uncertainty principle is true: the values of incompatible observables are not simultaneously verifiable. The point, here, is to question *ontological* versions of the principle.

4.4.1 Incompleteness

In particular, the point is to call into question the *completeness* of quantum theory. Recall (§1.2.3.9) the standard interpretation of quantum theory, according to which an observable, F , has a value for a system in a state, W , just in case W assigns probability 1 to some possible value of F (and 0 to the others). This interpretation immediately implies that there are no quantum states that assign simultaneously definite values to all observables, or indeed to any two observables with no common eigenvectors. Hence any argument that successfully shows that two such observables must have definite values implies that quantum theory is incomplete — more precisely, it shows that under the standard interpretation of the theory, quantum states do not describe (and cannot describe) the complete physical state of a system. The EPR argument purports to show exactly this claim.

4.4.2 The Generic Experiment

Generically, the experiment considered by EPR involves a pair of particles (call them α and β) prepared in the state

$$(93) \quad |\Psi_{\text{EPR}}\rangle = \frac{1}{\sqrt{N}} \sum_{n=1}^N |a_n\rangle |b_n\rangle$$

⁵⁹There are very good reasons to believe that the argument as presented by EPR [1935] was not quite what Einstein himself had intended. (The paper was not written by him.) See Fine [1986, esp. chs. 3-5], for example.

where the $|a_n\rangle$ and the $|b_n\rangle$ form orthonormal sets.⁶⁰ Hence there are observables, A for system α and B for system β (whose eigenvectors are respectively the $|a_n\rangle$ and the $|b_n\rangle$, corresponding to eigenvalues a_n and b_n) that are perfectly correlated (§1.2.6.4) in this state.

So suppose that α and β are in the state $|\Psi_{\text{EPR}}\rangle$ while they are spatially separated. Then the perfect correlation between A and B allows one to discover the value of A without, as EPR say, ‘in any way disturbing’ α , simply by measuring B on β .

Now for the crucial point. Consider orthonormal bases $\{|a'_n\rangle\}$ and $\{|b'_n\rangle\}$ obtained from the $|a_n\rangle$ and the $|b_n\rangle$ as follows:

$$(94) \quad |a'_n\rangle = \frac{1}{\sqrt{2}}(|a_n\rangle - i|a_{n+1}\rangle) \quad |b'_n\rangle = \frac{1}{\sqrt{2}}(|b_n\rangle + i|b_{n+1}\rangle),$$

where the sum in subscripted ‘ $n + 1$ ’ is modulo- N (i.e., $N + 1 = 1$). In this basis, the state $|\Psi_{\text{EPR}}\rangle$ has exactly the same form:

$$(95) \quad |\Psi_{\text{EPR}}\rangle = \frac{1}{\sqrt{N}} \sum_{n=1}^N |a'_n\rangle |b'_n\rangle.$$

(To verify, plug (94) into (95) and simplify. The ‘cross’ terms of the form $-i|a_n\rangle|b_{n+1}\rangle$ and $i|a_n\rangle|b_{n+1}\rangle$ cancel.) Hence there are additional observables, A' and B' (whose eigenvectors are respectively the $|a'_n\rangle$ and the $|b'_n\rangle$) that are *also* perfectly correlated in the state $|\Psi_{\text{EPR}}\rangle$. Moreover, A' does not commute with (is incompatible with) A — indeed, they share no eigenvectors — and similarly for B and B' . Again, we can discover the value of A' on α by measuring B' on β .⁶¹

Finally, notice that this entire description is *quantum-mechanical*. One sometimes hears the EPR experiment described in more or less classical terms. For example, letting A and B be position and momentum (see §4.4.3), one might be tempted to describe the preparation of the EPR state as follows: fire two particles of equal mass from a common source with equal (in magnitude) but oppositely directed forces. Their positions (distance from the source) and momenta (relative to the source) will then be perfectly correlated. This picture *is* tempting, but it is also completely wrong. Indeed, standard quantum mechanics implies that the state just described cannot be prepared in a way that would allow one to infer the position (or momentum) of α from that of β , because such inferences would require us to know, with precision, the position and momentum of the source, and such knowledge already violates the (epistemic) uncertainty principle.

4.4.3 Position and Momentum

Above we assumed that the perfectly correlated observables have a discrete spectrum. In fact, when they consider an explicit example (rather than the abstract

⁶⁰In fact, there is a ‘continuous’ version of (93), which is appropriate when we are considering observables, such as position and momentum, with continuous spectra. See §4.4.3.

⁶¹The situation as described here is not mathematically exactly the same as that considered by EPR, but the result is the same.

case as considered above), EPR refer to position and momentum, and a state that is perfectly correlated in position and momentum (in the sense that it is a simultaneous eigenstate of the sum of the momenta of the two particles, and the difference of their positions). Explicitly,

$$(96) \quad \Psi_{\text{EPR}}(x_1, x_2) = \int_{-\infty}^{+\infty} e^{(2\pi i/\hbar)(x_1 - x_2 + x_0)p} dp,$$

for some fixed x_0 . However, it is worth pointing out that in fact this state is not an allowed state of the system at all — it is not a vector in $L^2(\mathbb{R}^2)$. Moreover, this state, even if it could be prepared, necessarily spreads under any time evolution (associated with a finite potential energy), hence immediately becoming a state of *less* than perfect correlation.⁶²

To overcome these limitations of EPR's example, one could consider, for example, a narrow Gaussian that is very close to being a state of perfect correlation, but doing so makes the argument to come messy at best. And we should keep in mind that the discussion in terms of position and momentum is in fact only an *example* that EPR give to illustrate the main point.

In the end, in fact, it is easier to consider an example that involves bounded observables. The simplest case is the so-called 'singlet' state of a pair of spin-1/2 particles. Generically, this case corresponds (for an appropriate choice of the $|a_n\rangle$ and the $|b_n\rangle$) to the state $|\Psi_{\text{EPR}}\rangle$ above, for $N = 2$. We will just continue to speak of the observables A , A' , B , and B' , with the understanding that the perfect correlations discussed above obtain.

4.4.4 The Argument

How can this experimental situation be used to generate an argument for the incompleteness of quantum theory? The conclusion that EPR hope to establish is that A and A' both have a definite value simultaneously. Because they share no common eigenvectors, this conclusion is inconsistent with the standard interpretation of quantum states (see §1.2.3.9).

One tempting path to this conclusion involves presuming that once B has been measured on β , thereby establishing the value of A on α , we can measure A' on α and thereby establish its value directly. However, the problem with this suggestion should be clear: the disturbance theory of uncertainty can quickly be applied to conclude that the measurement of A' on α disturbs the previously established value of A . We are thus reminded that the point of this discussion is, in fact, to establish the definiteness of A and A' without 'in any way' disturbing α .

EPR's strategy involves two assumptions. The first, with which at least some versions of the standard interpretation can easily agree, is their 'criterion for physical reality', which asserts that whenever the value of an observable can be predicted with certainty, the observable actually has that value. (Notice that this criterion

⁶²By the state 'spreading', here, we mean, roughly, that it gets closer to a uniform distribution over \mathbb{R}^2 . See Dickson [2002b] for a discussion of these points, and further references.

is inconsistent with the sort of verificationist or operationalist views, mentioned above (§4.3.2), according to which a system has a value only if that value has been obtained as the result of a measurement.) Before we introduce the second assumption, let us see how far we can get with just this one.

Notice that if we measure B on β , we can predict the value of A on α with certainty. Similarly for B' and A' . Of course, we cannot measure *both* B and B' . Hence, instead of considering *actual* measurements of B and B' , let us consider non-actual, but possible (i.e., ‘counterfactual’) measurements of B and B' . We have:

Premise 1 *Possibly*, B is measured on β , and in this case, α has a definite value for A .

Premise 2 *Possibly*, B' is measured on β , and in this case, α has a definite value for A' .

From these two premises (which follow from the criterion for physical reality), EPR hope to conclude:

Conclusion *Possibly*, α has a definite value for both A and A' .

However, the Conclusion does not follow from the Premises 1 and 2. Indeed, the logical problem is, in part, that there is no guarantee that the possible conditions (‘ B is measured’ and ‘ B' is measured’) are co-possible. Indeed, as we know, they are not.⁶³

Hence EPR need another premise. They introduce a notion of ‘non-disturbance’ that is supposed to help patch up the argument: although the conditions (measurement of B) under which we can infer the definiteness of A on α are incompatible with the conditions (measurement of B') under which we can infer the definiteness of A' , the difference between them is supposed to make no difference to α , because they only involve a change of circumstances for β , which may be spatially separated from α .

However, not just any such principle will work. Consider, for example, the following:

Weak non-disturbance. If B is measured on β and (therefore, by the criterion for physical reality) A is definite for α , then: had we not measured B on β , α would still have had a definite value for A (and likewise, substituting primed observables for the unprimed ones).

This principle, which might be taken to deny that measuring B is what brings it about that A has a definite value, is insufficient to get EPR’s conclusion. They need, instead:

⁶³Consider the following analogous argument. (i) It is possible that the paper is burned, and in this case it will be reduced to ashes. (ii) It is possible that the paper is not burned, and in this case it will remain whole. Therefore, (iii) it is possible that the paper is both whole and reduced to ashes. Of course, this argument is invalid.

Strong non-disturbance. If B is measured on β and (therefore, by the criterion for physical reality) A is definite for α , then: had we *instead* measured B' on β , α would still have had a definite value for A (and likewise, substituting primed observables for the unprimed ones).

The weak principle is insufficient to get the conclusion because the introduction of the measurement of B' (as opposed to the mere absence of the measurement of B) could destroy essential features of the situation, and in particular features that permit the inference of properties of α (its value for A) from the results of measurements (of B) on β .⁶⁴

The principle of strong non-disturbance is supposed, of course, to capture some notion of ‘locality’. In particular, the idea is supposed to be that *nothing* that is done to β can have any affect on the properties of α , under the assumption that the two particles are space-like separated. Einstein’s theory of special relativity is presumably supposed to license this assumption.⁶⁵ In any case, with the principle of strong non-disturbance, EPR have a logically valid argument for the Conclusion, above. Indeed, they could (and seem to claim to) establish a stronger form of the conclusion, replacing ‘Possibly’ with ‘Actually’, as follows. Argue, further, that a *measurement* of B (or B') does not bring it about that α has a value for A (A'), so that α must have values for both A and A' even when B (B') is not measured.

4.4.5 Replies to EPR

One can, in fact, deny strong non-disturbance by denying locality. We shall consider the status of locality in quantum theory below (§6). Here, we consider two other replies to EPR.

We have already more or less encountered one of these replies: note that the EPR argument has no impact on those verificationist or operationalist views according to which an explicit measurement of a physical quantity is required not merely for us to know its value, but also for it to *have* a value. Of course, such views must deny the criterion for physical reality (which, keep in mind, is only a sufficient, not a necessary condition), and many find this principle quite compelling. (Apart from actually having the value in question, what else could, reasonably, make it be the case that we can predict its value with certainty?)

But Bohr offered a reply that appears to rely neither on this verificationist or operationalist strategy, nor — so he claims — on an explicit endorsement of non-locality. In particular, recall that EPR were aiming to avoid the ‘disturbance’ account of uncertainty, according to which, for example, a measurement of A'

⁶⁴The possible-worlds semantics for counterfactuals makes the point easy to see: while the closest ‘ B is not measured’-worlds to the ‘ B is measured and A is definite for α ’-worlds might all be ‘ A is definite for α ’-worlds, those closest worlds might not contain any ‘ B is measured’-worlds, so that the closest ‘ B is not measured but B' is’-worlds to the ‘ B is measured and A is definite for α ’-worlds need not be ‘ A is definite for α ’-worlds.

⁶⁵See Malament, this volume.

physically disturbs the value of A . Here, because we are only ever performing measurements on β , it is far from clear how a measurement of B' , for example, could disturb the value of A on α — after all, such a disturbance would have to be non-local. Bohr's reply denies strong non-disturbance, but without (so the claim goes) endorsing a physically direct (Bohr uses the term 'mechanical') disturbance of (e.g., the value of A for) α as a consequence of any measurement (e.g., of B) on β . Instead, recall the idea from above (§4.3.2), that the very well-definedness of certain physical quantities relies on certain physical conditions being in place.

Indeed, consider the EPR experiment, now, as Bohr (and EPR) did, in terms of position and momentum. We will assume (without loss of generality) that the positions and momenta are defined relative to the source. On Bohr's view, the well-definedness of the *sum* of the momenta for each particle (i.e., the total momentum for the system) is maintained just so long as the system (particles plus source) remains closed, i.e., just as long as total momentum is conserved. But a measurement of the position of β introduces a disturbance of β 's momentum. The system is no longer closed: β either loses momentum to or gains momentum from an object (the measuring apparatus) that is external to the system. But then the total momentum (relative to the source, which 'knows nothing' about this external influence on β) is no longer conserved, and thus the conditions for its well-definedness (relative to the source) are no longer in place, and therefore any inferences that we might have made about the well-definedness of α 's momentum, based on the well-definedness of the total momentum, are no longer valid. In other words, on Bohr's view, the conditions required (under the circumstances) for the well-definedness of α 's momentum are, when we are measuring β 's position, no longer in place. Note that the 'old' type of disturbance is still at work here — we are indeed supposing that β 's momentum is physically disturbed by a measurement of its position — but in addition there is another type of disturbance at work: the measurement of β 's position 'disturbs' the conditions required for the well-definedness of α 's momentum (under the circumstances). Thus Bohr believes that he can avoid EPR's conclusion by (in essence) denying strong non-disturbance, but in a way that does not entail a non-local disturbance, in the sense of an exchange of energy or momentum between β and α .

4.4.6 *Where We Stand*

Regardless, finally, of whether one finds any of these replies to EPR satisfactory (and we should not make any judgment on this point at least until we have considered the issues in §5 and §6), it should be emphasized that none of them shows that the EPR argument is unsound, much less invalid. Indeed, they are, in a sense, defensive maneuvers designed to articulate a view of quantum theory that avoids EPR's conclusion by denying, in a consistent and presumably plausible way, one of the premises of their argument. EPR's conclusion is thus far from called into question by such replies. Indeed, many philosophers of quantum theory are more or less in agreement with EPR that standard quantum theory is incomplete, if not

because of EPR's argument, then because of the so-called 'measurement problem', which we consider next.

5 THE 'MEASUREMENT PROBLEM'

The problem of measurement is, perhaps, the most discussed issue in the foundations of quantum theory, and has inspired numerous and varied interpretations of the theory, from the brilliant to the bizarre. In this section, I will first review the problem (§5.1), emphasizing its generality. Then I will consider some natural, but in the end unsatisfactory, reactions to the problem (§5.4). In the final section (§5.5), I will consider a few examples of interpretations of the theory, each of them largely directed towards solving the measurement problem.

5.1 *The Basic Problem*

5.1.1 *'Schrödinger's Cat'*

Recall (1.2.3.2.b) that the superposition principle implies that for any observable, F , if a system can have each (or indeed, just two of) the eigenvalues of F as its values for F , then it is also possible for a system to have (assuming the eigenstate-eigenvalue link, §1.2.3.9) *no* value for F (because it is in a superposition of the eigenstates of F).

If F is some unfamiliar observable of the unfamiliar quantum world, then perhaps one can live with this consequence. But what about observables on middle-sized solid objects that are the bread and butter of our everyday experience? What about 'the (approximate) location of the house', or 'the (approximate) momentum of the horse', and so on? In 1935, Schrödinger illustrated the sort of problem that we face with the following example.

One can even set up quite ridiculous cases. A cat is penned up in a steel chamber, along with the following device (which must be secured against direct interference by the cat): in a Geiger counter there is a tiny bit of radioactive substance, so small, that perhaps in the course of the hour one of the atoms decays, but also, with equal probability, perhaps none; if it happens, the counter tube discharges and through a relay releases a hammer which shatters a small flask of hydrocyanic acid. If one has left this entire system to itself for an hour, one would say that the cat still lives if meanwhile no atom has decayed. The psi-function of the entire system would express this by having in it the living and dead cat (pardon the expression) mixed or smeared out in equal parts. (Schrödinger [1935a])

The point, of course, is that after some stretch of time, the atom is in a superposition of 'decayed' and 'not decayed', and hence the hammer, poison, and ultimately

the cat, are in corresponding superpositions — in the case of the cat, a superposition of ‘alive’ and ‘dead’. But of course we never witness cats in such states. So apparently there is a serious problem with quantum theory.

5.2 Measurement

Notice that the cat serves, in essence, as a decay-indication device, a kind of crude measuring apparatus. Indeed, it is precisely Schrödinger’s point that one can magnify a superposition at the microscopic level (the atom) to a superposition at the macroscopic level (the cat), the level at which one is probably more inclined to deny categorically that superpositions (of at least some sorts, such as of the states ‘alive’ and ‘dead’) make any sense, physically, or at any rate that they exist with anything like the frequency that quantum theory apparently predicts. Measurement devices designed to measure quantum-mechanical observables are characteristically of this sort: they ‘magnify’ the state of some microscopic quantum system into the (indicator, or ‘pointer’) state of an apparatus that can be directly observed.⁶⁶

Of course, this magnification is a physical process, and as such, it is in principle to be described as a solution to Schrödinger’s equation (or some other quantum-theoretic equation of motion). Those equations are *linear*. (In fact, it is precisely this linearity that partially grounds the principle of superposition: any linear combination of solutions is again a solution.) Indeed, recall that the dynamical evolution of a system can be described in terms of a family of unitary operators, which are, of course, linear.

So let a given measurement-interaction (between some apparatus and some measured system) be described by the operator U . (As an example, recall §3.1.2.) Let the apparatus’ ‘ready-to-measure’ state be $|\Psi_0\rangle$ and let its ‘pointer’ states be $|\Psi_n\rangle$. For each of the eigenstates, $|f_n\rangle$ of the measured observable, F , we presume that the measurement-interaction results in an accurate indication of the result:

$$(97) \quad U(|f_n\rangle|\Psi_0\rangle) = |f_n\rangle|\Psi_n\rangle$$

so that the state $|\Psi_n\rangle$ indicates that the measured system has the value f_n for F . But then, by linearity,

$$(98) \quad U\left(\sum_n k_n |f_n\rangle|\Psi_0\rangle\right) = \sum_n k_n |f_n\rangle|\Psi_n\rangle.$$

Now we are in trouble. The eigenstate-eigenvalue link implies that this final state is one in which the pointer-observable for the apparatus in fact *has no value*.

⁶⁶Recall my earlier (§1.4.3) skepticism, however, about ‘pure’ states such as $|\text{cat dead}\rangle$. Does this skepticism suggest a way out of the measurement problem? Alas, it does not. We could just as well speak in terms of mixed states, here, and generate the problem, which relies only on the fact that at the end of a measurement, the state of the apparatus may assign non-trivial probabilities to *all* possible values of the ‘pointer-observable’.

The pointer is not indicating *anything*. But at the end of typical successful measurements, the apparatus *does* indicate a result, even when the measured system begins in a superposition of eigenstates of the measured observable. Standard quantum theory seems to contradict this apparent fact of experience. Hence the ‘measurement problem’.

5.3 *Generality of the Problem*

There are at least three senses in which the ‘measurement problem’ as described above does not capture the real problem with sufficient generality. First, it relies on the quite conservative eigenstate-eigenvalue link, and on a very restrictive notion of ‘measurement’. Second, it fails to indicate the ubiquity of superpositions. Third, it relies on the questionable assumption that the macroscopic pointer-states of the apparatus are pure states, when in fact they are almost surely mixed. In this subsection I will briefly consider these points. Finally, I will briefly raise a related problem, the problem of the ‘classical limit’.

5.3.1 *‘No-go’ Theorems*

Two important premises in the argument leading up to the measurement problem, above, were the eigenstate-eigenvalue link and the account of what counts as a successful ‘measurement’. Both of these assumptions can be weakened considerably.⁶⁷ I will consider each in turn.

5.3.1.1 Weaker Conditions for Definiteness First, we may introduce a weak condition for when an observable has a definite value, roughly as follows: the final state of the apparatus assigns a definite value to the pointer-observable for the apparatus just in case the final state of the compound system is a mixture of states each of which has a definite value for the pointer observable according to the eigenstate-eigenvalue link. In other words, we now allow the adoption of an ignorance interpretation of mixed states of the compound system, accepting that when the compound system has the sort of mixed state just described, it is actually in one of the pure states appearing in the mixture, and therefore, according to the eigenstate-eigenvalue link, the apparatus has a definite value for the pointer-observable.

Note that the condition applies to the compound system, not to the apparatus on its own. The difference is subtle, but crucial. For example, in an ideal measurement, the apparatus by itself will always be in a mixture of eigenstates of the pointer-observable. It does not follow, however, that this mixture can be given an ignorance interpretation. I shall discuss this point in a somewhat different context below (§5.4.5.2.a).

⁶⁷There is a long history of proofs, increasingly general, that the measurement problem is ‘insoluble’ in the context of standard quantum theory. A review of the early history, together with arguably the simplest proof of the theorem, is given by Brown [1986]. For a more recent survey and extensive discussion of this issue, see Mittestaedt [1998, esp. ch. 4].

5.3.1.2 Weaker Accounts of Measurement We may also weaken the account of measurement, along the following lines. Let U represent the time-evolution of states of the compound system during the measurement of a POVM, E , by a ‘pointer-observable’ POVM, Q . We require: for any two states, W and W' , of the compound system, if W and W' differ in their probabilities for at least one of the effects in the image of E , then UWU^{-1} and $UW'U^{-1}$ differ in their probabilities for at least one of the effects in the image of Q . Intuitively, U renders the pointer-observable ‘somehow sensitive (even if only probabilistically) to the measured observable’.

Note, however, that while this weaker account of measurement, together with the weaker criterion for definiteness (§5.3.1.1), is sufficient to derive the measurement problem (i.e., the non-definiteness of the pointer-observable at the end of a measurement), it is also worth keeping in mind that, by the biorthogonal decomposition theorem (§1.2.6.2), there are *some* observables (one for the measured system and one for the apparatus) with respect to which the state of the compound system has the same form that it does in an ideal measurement. (I.e., write the state in its biorthogonal form. The bases for the measured system and apparatus will be bases determined by *some* observable for each.) On the other hand, whether *this* apparatus-observable is one that we antecedently believe to be definite is another question, the answer to which probably depends on the details of the interaction.

5.3.2 The Ubiquity of Superposition

The superposition principle asserts that superpositions of possible states are again possible states. But perhaps they are extremely rare. In that case, again perhaps we should not be terribly concerned about the measurement problem.

In fact, however, superpositions at the microscopic level are ubiquitous. To see why, we need only consider the case of spin. Recall that $S_{\vec{u}}$ and $S_{\vec{u}'}$ are incompatible unless $\vec{u} = \vec{u}'$ or $\vec{u} = -\vec{u}'$. Hence *every* spin-1/2 particle is ‘in a superposition’ with respect to just about every direction of spin. A similar remark holds for the polarization of photons. And there is good experimental and theoretical reason to believe that many fundamental particles may have wavefunctions that are highly de-localized.

Moreover, we need not imagine exotic situations such as that described by Schrödinger in order to believe that these microscopic superpositions may be ‘magnified’. While such magnification is an important part of measurement, there is little reason to suppose that it does not occur naturally, and frequently. After all, the human eye, for example, is sensitive to as few as a half-dozen or so photons. It seems plausible to suppose that many other interactions that occur in nature have the effect of correlating the state of some macroscopic object with the state of some microscopic object — and such correlation (even if imperfect) is sufficient for the sort of magnification that will give rise to the measurement problem.

Hence, there is good reason to believe that the measurement problem is highly general, in a few senses: (a) even under apparently weak accounts of when appa-

ratus' have a definite value for some macroscopic observable (§5.3.1.1), quantum theory apparently implies that they do not; (b) the sort of correlation between a macroscopic and microscopic object that is required in order to bring about the situation in (a) is very weak, and moreover perfect correlations are, by the biorthogonal decomposition theorem, ubiquitous (§5.3.1.2); and (c) the sort of superposition, and interaction, that is required give rise to the situation in (b) is plausibly quite common (this section).

5.3.3 *The Classical Limit*

Finally, I will quickly notice a problem that is discussed in detail elsewhere in this volume. In a sense, the measurement problem is the problem that the 'weirdness' of the quantum world is somehow not confined to the microscopic world. In other words, the 'weird' quantum world does not (always) mesh nicely with our more familiar classical world, in the sense that superpositions at the micro-level (where we might be able to live with them) are not confined to that level, but can be made to appear at the macro-level (where, so the argument goes, we cannot live with them). This apparent conflict leads to a more general question: how does quantum theory go over to classical theory (which, after all, works quite well for a large domain of objects in a wide variety of conditions)? The question is vexed in a number of ways, which, however, I shall not consider here. I refer the reader to Landsman, Ch. 5, this volume (esp. §§5,6).

5.4 *Non-Solutions*

Numerous solutions to the measurement problem have been proposed, and we shall consider some of them in a separate section on interpretations. Here, we consider several proposed solutions that in fact do not work, or at the very least, face extraordinary obstacles to making them work.

5.4.1 *Naïve Realism*

5.4.1.1 The Obvious Solution to the Rescue? The obvious solution to the problem is to give up the eigenstate-eigenvalue link. Indeed, why not interpret the probabilities delivered by quantum theory as entirely epistemic? That is, why not suppose that *every* observable has a definite value *all of the time*, and that incompatibility merely represents the fact that observing the value of one of them disturbs the values of others in an uncontrollable way? Something like an epistemic version of Bohr's approach to uncertainty, if it works, could be put to use in the service of this view, both accounting for the uncertainty relations, and avoiding the measurement problem.

5.4.1.2 The Kochen-Specker Theorem While initially attractive, this view founders on a well-known theorem, the Kochen-Specker theorem (Kochen and Specker [1967]). Notice that naïve realism seeks to identify a single value for each

observable. With one further requirement, one can show that no such assignment is possible.

5.4.1.2.a Non-Contextuality This further requirement is specifiable in a variety of ways. Here we put it in terms of ‘non-contextuality’. Note that assigning a value to an observable amounts to assigning, to each set of mutually orthogonal subspaces that span the entire space, a ‘1’ to exactly one of the subspaces, and ‘0’ to the others. (These subspaces are the eigenspaces of the observable to which a value is assigned; assigning a 1 to a given eigensubspace is equivalent to assigning the corresponding eigenvalue to the observable.⁶⁸) Proceeding in this way, we will further assume that all observables with exactly the same eigenspaces have ‘the same’ value, in the sense that they have the eigenvalue associated with a given common eigenspace. This assumption amounts to the requirement that for any observable G , if G has the value g , then for any function, $f()$, the observable $f(G)$ has the value $f(g)$.⁶⁹

Non-contextuality includes this assumption, but goes a step further, requiring, in addition, that the assignment of a ‘0’ or ‘1’ to a given subspace P is independent of *which* set of mutually orthogonal and jointly spanning subspaces P is considered to be a member. This requirement amounts to the condition that whenever two observables, F and G , share an eigenspace, P (but do not necessarily share *all* of their eigenspaces), F has the eigenvalue corresponding to P if and only if G does. (The eigenvalues will in general differ, of course.) (Note that in this case F and G in general will not commute.)

5.4.1.2.b The State-Independent Theorem Given non-contextuality, naïve realism amounts to the requirement that we be able to find a homomorphism from the lattice, $L(\mathcal{H})$, of subspaces of a Hilbert space, \mathcal{H} , to the Boolean lattice $\{0, 1\}$. That is, we require a map, $h : L(\mathcal{H}) \rightarrow \{0, 1\}$ such that, for any $P, Q \in L(\mathcal{H})$, $h(P) \leq h(Q)$ if and only if $P \leq Q$. It follows (from the definitions of the meet, join, and orthocomplement, which are all given in terms of the partial order — §7.4.2), that $h(P) \wedge h(Q) = h(P \wedge Q)$, $h(P) \vee h(Q) = h(P \vee Q)$, and $h(P)^\perp = h(P^\perp)$. But notice that in $\{0, 1\}$, the operators \wedge , \vee , and $^\perp$ behave just as those from classical logic (i.e.: $0 \wedge 0 = 0$, $0 \wedge 1 = 0$, $1 \wedge 1 = 1$, $0 \vee 0 = 0$, $0 \vee 1 = 1$, $1 \vee 1 = 1$, and $0^\perp = 1$). In other words, the naïve realist requires that it be possible to consider quantum theory to arise from an underlying (logically) ‘classical’ theory.

Such a picture *is* available in classical mechanics. Indeed, let Γ be the phase space for a classical system. The physical propositions (cf. §2.1.1) are represented by the lattice of Borel subsets of Γ (where the partial order is given by subset inclusion), and it *is* possible to define a homomorphism from this algebra to $\{0, 1\}$.

⁶⁸Here I ignore observables, such as position and momentum, whose spectrum is continuous (so that their possible values do not correspond to subspaces). After all, assigning them values *as well* only makes the naïve realist’s life even harder.

⁶⁹Note that, in general, if $G = \sum_n g_n P_n$ is the spectral decomposition of G , then $f(G) = \sum_n f(g_n) P_n$.

Indeed, choose a point $x \in \Gamma$, and define the map, δ_x , by, for any Borel subset $S \subseteq \Gamma$, $\delta_x(S) = 1$ if $x \in S$ and 0 otherwise. This δ_x is a homomorphism.

The content of the Kochen-Specker theorem is that for lattices of subspaces of a Hilbert space whose dimension is greater than 2, there are no such homomorphisms. This version of the theorem is sometimes called ‘state-independent’ because it does not rely on any presumptions about the quantum state of a system, but only on the *structure* of the state space as a whole.

5.4.1.2.c State-Dependent Theorems There are other versions of the theorem that begin with the assumption that the system is in some given quantum state. In this case, we can introduce another condition on the value assignments: they must respect the probabilities generated by the quantum state. For example, if the state assigns probability 0 [1] to a given value, f , for F , then the value-assignment must assign 0 [1] to the corresponding eigenspace. There are some particularly simple state-dependent Kochen-Specker theorems, the most famous being the GHZ (Greenberger-Horne-Zeilinger [Greenberger *et al.*, 1989]) theorem, which (in a form nicely described by Mermin [1990]) considers a three-particle system of spin-1/2 particles, and the observables

$$(99) \quad S_x^1, S_y^1, S_x^2, S_y^2, S_x^3, S_y^3,$$

$$(100) \quad S_x^1 \otimes S_y^2 \otimes S_x^3, \quad S_y^1 \otimes S_x^2 \otimes S_y^3, \quad S_y^1 \otimes S_y^2 \otimes S_x^3, \text{ and } S_x^1 \otimes S_x^2 \otimes S_x^3,$$

where S_x^1 is shorthand for $S_x^1 \otimes \mathbb{I}^2 \otimes \mathbb{I}^3$ and the superscripts indicate the particles with which each observable is associated. For simplicity, suppose that each of the $S_u^{(n)}$ has been ‘normalized’ to have eigenvalues ± 1 (rather than the usual $\pm 1/2$). Hence the eigenvalues of all of the observables in (99) and (100) are ± 1 .

Note that the observables in (100) commute. Hence we can consider a state, the ‘GHZ-state’, that is a simultaneous eigenstate of them all. We shall consider an eigenstate with eigenvalues $+1, +1, +1, -1$ respectively. For a system in this quantum state, any ascription of definite values to the observables *must* assign those values to the observables in (100). Let $v()$ be a map from the observables in (99) and (100) to their values, and consider then the following array of values:

$$(101) \quad \begin{array}{ccc} v(S_x^1) & v(S_y^2) & v(S_y^3) \\ v(S_y^1) & v(S_x^2) & v(S_y^3) \\ v(S_y^1) & v(S_y^2) & v(S_x^3) \\ v(S_x^1) & v(S_x^2) & v(S_x^3). \end{array}$$

The possible values are always ± 1 . The product of the values *across* each row must be $+1$ for the first three rows, and -1 for the last, so that the product of all twelve numbers must be -1 . But on the other hand, $v(S_u^n)$ for each $n = 1, 2, 3$ and $u = x, y$ appears exactly twice in the array, so that the product of all twelve numbers must be $+1$, contradicting our conclusion above. Hence there is no assignment of values to all ten of the observables in (99) and (100) that is consistent with the GHZ-state.

Naïve realism is thus in serious trouble. The most straightforward reading of its basic commitments leads, via the Kochen-Specker theorem and its analogues such as the GHZ theorem, to a logical contradiction. While there are proposals to save aspects of naïve realism from this problem, all of them (of course) violate one or more of the conditions of the Kochen-Specker theorem, and, arguably, are no longer ‘naïve’ realism.

5.4.2 Ensemble Interpretations

Ensemble interpretations (e.g., Ballentine [1970]) attempt to avoid the measurement problem altogether by stipulating that quantum states are not *about* individual quantum systems. They are essentially statistical, and thus can be used to describe only ensembles of systems. Hence the state at the end of a measurement ought not be thought of as describing a particular apparatus (and measured system) at the end of a particular measurement, but all such apparatuses (and measured systems) and the end of all such measurements.

There are two strains of this sort of interpretation in the literature, not always distinguished. The first we might call ‘minimalist’: it is the deflationary view that one can make *sense* of the quantum state (for example, at the end of a measurement) only as the description of an ensemble of similarly prepared systems.

The second strain of the ensemble interpretation makes a stronger claim, namely, something to the following effect: nothing more than the deflationary account of the quantum state is needed; i.e., there is no scientific demand for a theory that describes individual systems. On this view, such a demand is an artifact of a deterministic world view, and the usual ‘mysteries’ of quantum theory simply disappear once one learns to live with statistical theories. However, most researchers in the philosophy and foundations of quantum theory appear to be unconvinced.

Moreover, even the minimalist claim faces a problem, and the problem also shows why the stronger claims are, at the least, in need of further development. The Kochen-Specker theorem already shows that the probabilities generated by quantum states cannot be understood straightforwardly as classical probabilities. Thus the ensemble interpretations appear to be pushed towards other views, such as the quantum logic view, which will be considered later.

5.4.3 Collapse Postulate

The ‘standard’ solution to the problem of measurement is the so-called ‘collapse postulate’ (also called the ‘projection postulate’) of quantum theory, already mentioned in §3.1.2, first discussed in detail by von Neumann ([1932, p. 351 and pp. 417-418 of the English translation]). The postulate can be found stated in numerous ways, such as:

Collapse Postulate/ Upon measurement of the observable F on a system in the state W , the result of the measurement will be an eigenvalue corresponding to some eigenspace, P of F , and the state of the

system will then be $PWP/\text{Tr}[PWP]$.

If the state can be written as a vector, $|\psi\rangle$, then the ‘collapse’ amounts to projecting this vector onto $\text{ran } P$, i.e., onto $P|\psi\rangle$, and renormalizing the result.

There are many reasons to be unhappy with the collapse postulate, but it is worth saying, at the start, that most of the time, for most of the predictions that one wants to make from quantum theory, the postulate *works* very well.

One often-cited problem with the collapse postulate is that it amounts to discontinuous (and irreversible) evolution of a system. The resulting picture of the evolution of quantum systems is thus odd indeed: continuous, deterministic, reversible, unitary evolution (through Hilbert space) according to the Schrödinger equation (or some other quantum equation of motion), punctuated by discontinuous, irreversible, probabilistic, instantaneous, changes of the state.

At the very least, one would like some story about why these punctuations occur. Clearly, in the statement of the postulate, they are connected with measurement. But what is physically special about measurement? Alas, no compelling answer is forthcoming. Indeed, the most obvious problem with the collapse postulate is that it relies on an unanalyzed notion of ‘measurement’. Most of the time, reasonable people can agree about when a measurement has occurred, but that fact is entirely unhelpful. For most researchers in the foundations and philosophy of quantum theory seem to agree that if quantum theory is supposed to be a fundamental theory, then it should tell *us* when a measurement occurs (or more generally, what a measurement is), not vice versa. The world should not rely on *us* to ‘tell it’ when to collapse a state.

5.4.4 Macro-Micro Distinction

One answer to the question ‘what is a measurement’ is suggested by our discussion of Schrödinger’s cat (§5.1.1), and is sometimes put forward as a response to the measurement problem. The suggestion is that a measurement occurs when the state of a microscopic system is ‘magnified’ in such a way that it becomes correlated with the state of some macroscopic system. There is no denying that paradigmatic measurements do have this feature. Alas, this characterization relies on another ill-defined notion, namely, the difference between the microscopic and the macroscopic.

The distinction is also sometimes made between the ‘easily reversible’ and the ‘essentially irreversible’. Because microscopic systems have few degrees of freedom, their behavior is often easy to reverse, while the behavior of macroscopic systems, with vastly many degrees of freedom, is very difficult if not practically impossible to reverse.

In either case — whether the point is made in terms of size or reversibility — the idea is supposed to be that at an appropriately large (or irreversible) scale, physical systems behave classically. Indeed, one sometimes hears ‘the Copenhagen’⁷⁰

⁷⁰This ‘interpretation’ — unlikely, however, to have ever been a single unified view — is

interpretation of quantum theory, and especially Bohr's version of it, characterized in this way: one *must* presume that measuring apparatuses are classical objects, and this presumption provides the grounds for determining when a collapse (measurement) occurs.

However, while again providing an often useful practical characterization of measurement and 'the classical', these distinctions are, in most researchers' view, not well enough defined to sit at the foundations of a supposedly fundamental theory.

5.4.5 Decoherence

A related idea, but recently more well developed, is that sufficient interaction with 'the environment' serves, in essence, to 'collapse' ('decohere') the state of a system. Here we need to be very careful, because there is a well-studied phenomenon, called 'decoherence', that involves the interaction of a system with its environment. There is no disputing the physical importance of this phenomenon. On the other hand, many have invoked this phenomenon as providing some sort of solution to the measurement problem. This claim requires careful scrutiny.

5.4.5.1 The Phenomenon of Decoherence

5.4.5.1.a Qualitative Description It is remarkably difficult to isolate a physical system from the rest of the world (its 'environment'). Particles from almost every corner of the particle zoo are hurling around, and a great many of them can penetrate even very strong barriers (such as lead walls). Even for very small systems (such as a particle of dust), it is well-nigh impossible to prevent significant interaction with the environment.

Physicists have developed both simple and very sophisticated models of this interaction, ranging from assuming that the system interacts with an otherwise unspecified thermal bath to carefully modeling the rate, nature, and strength of interactions that a given system is likely to experience in a given type of environment. With these models, one can estimate (and in very rare, usually highly idealized, cases, explicitly determine) the effect of these interactions on the state of the system.

Notice that these interactions will *entangle* the system with the environment. We are therefore no longer talking, in the first place, about 'the' state of the system, but about the state of the composite 'system-plus-environment'. Of course, typically we will have little or no access to the relevant environmental degrees of freedom. (Imagine, for example, that a photon bounces off a dust particle, then gets trapped in the atmosphere, or worse, heads out to space. In the former case, recovering the photon will be practically impossible, and in the latter case, it could

so-called because it is typically associated with Niels Bohr and his associates, who worked in Copenhagen. See Cushing [1994, chs. 6,7], Beller [1999], and references therein, as well as Landsman, Ch. 4, §1, this volume.

be literally impossible.) Although these degrees of freedom in a sense ‘encode’ certain information about the system (such as its location, because typically the interactions are position-dependent), that information is almost always practically, and occasionally in principle, lost to us.

Hence, although the system will become entangled with the environment, we typically have access only to its reduced state, not to the state of the total composite. We obtain this reduced state by tracing out the degrees of freedom of the environment. In many models of the interaction, the result is a reduced state for the system that is approximately diagonal in position; that is, its (mixed) state looks like the state of a system that is well-localized in space (in the sense that each component (spectral projection) of the mixture is well-localized in space). Decoherence may, in this sense, be thought of as ‘localizing’ the system (but see section 5.4.5.2.a, where we will worry explicitly about the fact that this mixture is improper).

This localization ultimately derives from the fact that in these models, the interaction between the system and the environment is position-dependent. (The environment interacts with the particle only in the vicinity of the particle.) More generally, if the interaction Hamiltonian that describes the exchanges of energy between the system and its environment commutes with some system observable, $Q \otimes \mathbb{I}$, then the reduced state of the system becomes approximately diagonal in the basis picked out by the eigenvectors (or eigenspaces) of Q . The environment is said to ‘suppress’ the off-diagonal, ‘interference’ terms. Moreover, because the interaction Hamiltonian commutes with Q , a system that is already in an eigenstate of Q will tend to remain so (assuming, as is often physically reasonable, that the interaction Hamiltonian swamps the effects of the free Hamiltonian for the system).

5.4.5.1.b Example A classic and much-studied example of decoherence (e.g., Joos and Zeh [1985]), involves a dust particle in the atmosphere. Here is a simplified summary of their argument.

Let $|\psi\rangle$ represent the initial state of the dust particle. Let the states $|\psi_q\rangle$ be a basis of well-localized states for the particle (each centered at a position labeled q). Let $|E_0\rangle$ be the initial state of the environment, and consider the interaction between the particle and a single air molecule in the environment for a particle that is in one of the states $|\psi_q\rangle$: as a result of this single interaction, we will assume that the particle-plus-environment evolves into the state $|\psi_q\rangle \otimes |E_q\rangle$. (Here we assume, in particular, that the dust particle is much heavier than the air molecule, so that the interaction leaves the state of the dust molecule essentially unchanged). Joos and Zeh show, roughly, that if $|\psi\rangle$ (represented as a wavefunction) is initially in a Gaussian (not necessarily well-localized) state, written in the $|\psi_q\rangle$ -basis as $|\psi\rangle = \sum_q c_q |\psi_q\rangle$, then the matrix representation of the reduced density operator for the particle in the $|\psi_q\rangle$ -basis (recall §1.2.3.4), after one such interaction, will be:

$$(102) \quad W_{qq'} = c_q c_q^* \langle E_q | E_{q'} \rangle,$$

where $|\langle E_q | E_{q'} \rangle| \approx 0$ whenever the distance between q and q' is much larger than the wavelength of the dust particle.⁷¹ In other words, the off-diagonal ($q \neq q'$) terms in $W_{qq'}$ get reduced by a factor of $\langle E_q | E_{q'} \rangle$. The intuition here is that because the states $|E_q\rangle$ and $|E_{q'}\rangle$ correspond to environments where the air molecule has scattered from the dust particle in two different locations (represented by q and q'), if those locations are very different, then the corresponding states of the environment will be ‘very different’ (i.e., nearly orthogonal). (Of course, if $q = q'$ then this inner product is 1.)

Joos and Zeh then show, for a wide range of different models of the environment, that after *many* such interactions, the off-diagonal terms in the reduced density matrix decay exponentially, at a rate that depends on $q - q'$ (and is zero when $q = q'$). The rate is very fast: according to Joos [1986], the reduced density matrix for a dust particle with radius 10^{-5} cm in even a high quality vacuum will be extremely close to diagonal in position⁷² in about one microsecond. One says that the state of the object system (the dust particle) has ‘decohered’ as a result of its interaction with the environment.

5.4.5.2 Decoherence and the Measurement Problem

5.4.5.2.a Decoherence Does not Solve the Problem Decoherence appears to ‘localize’ systems that interact with their environment, so that perhaps we will not, after all, ever need to worry about encountering cats that are ‘smeared’ between two macroscopically distinct states (recall §5.1.1). Alas, there is a problem with this conclusion, stemming from the fact that the reduced state of the system is an *improper* mixture (recall §1.2.6.3.b). Let us investigate this point further.

On the one hand, we must acknowledge that it would be practically impossible to perform an *experiment* to determine that the reduced state of a system that has decohered is *not* a proper mixture. To see why, consider the very simple case of two perfectly correlated particles (analogous to the object system and its environment, which, as a result of their interaction, have become correlated), as in (93). The reduced state for α is $W_\alpha = \sum_n (1/N) |a_n\rangle \langle a_n|$ (analogous to the reduced state of the object system, which, recall, is ‘nearly diagonal in position’). This state makes the same predictions for *every* observable on α as the state in (93) does. In order to distinguish between proper and improper W_α , we would therefore need to measure some observable on the compound system ($\alpha\&\beta$), and in particular we would need to measure the *correlations* between α and β (for α & β is *not* in an entangled state — and therefore α ’s state is a *proper* mixture — if and only if there are no non-trivial correlations between them). Now, in the case

⁷¹Here we are relying on de Broglie’s relation between momentum and wavelength — recall note 57 of §4.2.1.

⁷²In this case, ‘extremely close to diagonal in position’ means that the reduced density operator can be written as a sum of states that represent the dust particle localized to within about 10^{-13} cm.

of (93), measuring these correlations is relatively straightforward (assuming that we have many copies of the pair all in the same state), because in fact they are perfectly correlated in many observables. But in the case of interest here — an object system and its environment — things are decidedly more difficult, because the correlations exist only between a very few particles of the environment and the object system. As I noted above, it is in general practically impossible to recover those particles from the environment, and even if we could, in general not just any measurement on them will do the job. Indeed, if we do not have many copies of the total system (and in general we will not!), then the measurement in question becomes even more difficult, because we must measure an observable (analogous to the projection onto the state in eq. 93) that is not even a product (i.e., of the form $F \otimes G$).

On the other hand, the practical impossibility of performing such a measurement on the environment does not by itself license the assumption that the mixed state of the decohered system is (or may be treated as if it is) proper. The problem, as ever, is the eigenstate-eigenvalue link. The *true* quantum-mechanical state of the compound system is an entangled state, and in this state, according to the eigenstate-eigenvalue link, the object system does *not* have a definite location (in general). The assumption that the state of the system ‘might as well’ be a proper mixture is in fact *inconsistent* with this claim. In other words, while the two states are observationally very hard to distinguish, they are (given the eigenstate-eigenvalue link) interpretationally inconsistent.

Of course, one could give up on the eigenstate-eigenvalue link, and many interpretations do. We will discuss some of them shortly. The point here is that, with the eigenstate-eigenvalue link in place, decoherence does not solve the measurement problem.

5.4.5.2.b Decoherence Does Help with the Problem Decoherence does, however, help with the most general form of the problem of measurement. Recall (§5.3.2) that the measurement problem is not confined to interactions that would typically be regarded as measurements. *Many* interactions between quantum and macroscopic systems will put the macroscopic system in a state that, given the eigenstate-eigenvalue link, will be contrary to what most presume to be the deliverances of everyday experience. In other words, in its most general form, the ‘measurement’ problem is just the problem that quantum theory apparently fails to assign definite values to observables that appear, on the basis of everyday observation, to have definite values. In the special case of a typical measurement, we are perhaps already halfway to a solution, for in this case, the apparatus is at least already in a mixture, albeit improper, of ‘desired’ states.

But what about the more general sorts of interactions, mentioned above, that are not explicitly measurements but still generate ‘bad’ states for macroscopic objects? Decoherence promises to secure the following: for ‘relatively large’ systems (for example, at least as big as a dust particle), interaction with the environment will entail that the state of the system will become a mixture, albeit improper, of

‘desired’ states (or something that is very close to such a mixture — so close that one might be willing to overlook the difference). Even so, however, this strategy has merely reduced the general problem to a problem that continues to elude a satisfactory solution, although it does give one further justification for focusing the discussion specifically on measurements (because decoherence apparently reduces the general situation to one that is at least formally similar to measurement-situations).

5.5 Interpretations

5.5.1 Ways of Pursuing the Project

I have covered a number of philosophical and foundational issues arising from quantum theory, and an ‘interpretation’ of the theory should address all of them in some way or other — from providing an account of the empirical content of the theory, to providing an understanding of incompatibility (and thereby the ‘uncertainty relations’), to providing some understanding of the apparent failure of locality in the theory (§6). Nonetheless, most interpretations are aimed primarily at solving the measurement problem. As we saw above, that problem can be characterized (among other ways) as a conflict between (i) a common understanding of the physical properties of familiar physical objects, including the sorts of inferences about them that are valid, and (ii) a minimal account of when observables are empirically well-defined, taken in the context of (iii) quantum theory. Conceived in this way, there are three ingredients that lead to a contradiction, and therefore three generic strategies for avoiding the contradiction:

1. Deny that the common understanding of the physical world or inferences made about the physical world are always true or valid. For example, one might deny, contrary to appearance, that ‘pointers’ at the end of a measurement have a single definite (or nearly definite) position. We will encounter other ways to ‘deny common sense’ below. Let us call these theories ‘uncommon sense interpretations’.
2. Supplement, or replace, the minimal account of when observables are empirically well-defined. In some cases, extra ‘hidden’ variables are added to the theory. In all cases, systems are said to have properties that go beyond what standard quantum theory would assign, and hence all such interpretations are often, with admitted stretching of the concept, called ‘hidden-variables interpretations’. (Inasmuch as adding new physical variables — and often a dynamics for them — to the theory is tantamount to proposing a new theory, these interpretations are often called ‘hidden-variables *theories*’. I shall usually stick with the term ‘interpretation’, though I mean to be making no particularly substantive claim by doing so, and I shall occasionally use the term ‘theory’ to conform to custom.)
3. Supplement (i.e., change) standard quantum theory as thus far described

with some additional physics. Such interpretations add a rule (usually conceived as expressing a dynamical law of some sort) that results in systems having the properties, or nearly the properties, that we naïvely take them to have. Because the result of all such rules is a ‘collapse’ of the state from a superposition of values to a single value, these interpretations are often called ‘collapse interpretations’. (Again, they are sometimes instead called ‘collapse theories’ and I shall occasionally use that term as well.)

(Of course, the strategies might also be combined, but generally interpretations characterize themselves in terms of one of these strategies, even if one or both of the others fall out as an additional consequence.) I shall accordingly discuss interpretations under these three headings.

There is not space here to consider *all*, or even a *majority*, of the interpretations in each class. I will therefore have to restrict attention to a brief account of just two important representatives from each of the three genres. The reader should keep in mind that there are different versions of even the few interpretations that I shall discuss here. In most cases, I have chosen to describe a version that strikes me as the least problematic, but this claim is, in most cases, controversial.

My treatment of these interpretations will, necessarily, be brief. I will give a few more references than I have previously, to get the interested reader started on the literature.

5.5.2 ‘Uncommon Sense’ Interpretations

5.5.2.1 Quantum Logic Interpretations I have already noticed (§2.1.1) that the lattice of subspaces of a Hilbert space, i.e., the lattice of ‘propositions’ about a physical system, can, *prima facie*, be interpreted logically. Quantum logic interpretations take this idea seriously, and understand quantum theory to necessitate a revolutionary change from classical to quantum logic.⁷³

One way to characterize the difference between classical and quantum logic is in terms of the failure of distributivity (§7.4.3) — the classical and quantum lattices of propositions are otherwise structurally (logically) the same. The basic idea of the quantum logic interpretation, then, is that in particular the classical law of distributivity is *invalid*. This invalidity is supposed to allow one to make the sorts of assertion that, for example, the naïve realist (§5.4.1) had hoped all along to be able to make. For example, let F and G be two incompatible maximal observables sharing no eigenvectors. Denote their eigenspaces by $\{F_n\}$ and $\{G_m\}$. As propositions (elements of the lattice of subspaces of the Hilbert space, interpreted

⁷³ Birkhoff and von Neumann [1936] were among the earliest advocates of something like a quantum logic interpretation. Subsequent work in this area is founded on theirs. Some of the advocates of a quantum logic interpretation have been: Finklestein [1962; 1969], Putnam [1969], Friedman and Putnam [1978], and Bub [1974]. See also the collection of papers (Hooker [1975; 1979]). Some well-known critiques have been made by Dummett [1976] and Gibbons [1987]. Many of the standard objections are addressed in Dickson [2001], and occasionally also in the work of the advocates listed above.

logically), they assert that a system has the corresponding eigenvalue as a value for the observable. Then:

$$(103) \left(\bigvee_n F_n \right) \wedge \left(\bigvee_m G_m \right) = \mathbb{I} \wedge \mathbb{I} = \mathbb{I}$$

where \mathbb{I} is the logically true proposition. If we read the first half of this conjunction as asserting that the observable F has *some* value (understanding the existential quantifier in terms of disjunction, as one often does), and similarly for the second half of the conjunction, then (103) asserts that F has a value and G has a value. Indeed, this assertion is *logically* true. (Note that if we could apply distributivity to (103), then we could quickly turn it into a logical contradiction, because for any n and m , $F_n \wedge G_m = 0$, the logically false proposition.)

Moreover, the corresponding claim is true for a similar conjunction (\wedge) of such disjunctions for any number of observables. Hence (understanding the universal quantifier in terms of conjunction) the quantum logic interpretation claims to recover the idea that all observables (for a given system) always have a value (for that system). If so, then the measurement problem is no longer a problem. Of course, this approach also raises some questions. In addition to doubts about whether logic is revisable at all, some have argued that the quantum logic interpretation simply moves the mystery from one place (quantum theory) to another (logic). In any case, it is clear that the quantum logic interpretation needs to make some argument about why classical logic does in fact work within certain domains. (Decoherence could be helpful, here.) It also needs to say something about the success of classical reasoning in mathematics. Both of these projects amount to open questions in the quantum-logic interpretation (Dickson [2001]).

5.5.2.2 Many-Somethings Interpretations Our second example of ‘uncommon-sense’ interpretations is the ‘many-somethings’ interpretations. These interpretations generally go under the heading ‘many-worlds’ interpretations, but the notion of a ‘world’ that is at work here (quite problematic in its own right) is sufficiently far from the usual understanding of that term that a less committal word seems appropriate.⁷⁴

In any case, whereas quantum-logic interpretations deny (or at best re-interpret) the logical *validity* of apparently valid inferences about the properties of physical objects, many-somethings interpretations deny (or at best re-interpret) the *truth* of apparently true claims about the properties of physical objects.

Consider again the right-hand side of (98). The problem that this state raised earlier was that, on the standard interpretation (adopting the eigenstate-eigenvalue

⁷⁴Indeed, we are talking, here, about a rather wide class of interpretations. It is far from clear that the intention of the originator of this class of views, Everett [1957], was to invoke the idea of many worlds, which appears to have been introduced by DeWitt (see his, and Everett’s, article in DeWitt and Graham, [1973]). A ‘many-minds’ variant has been proposed by Albert and Loewer [1988] and Donald ([1990]), among others. See the review by Butterfield ([1995]) and the book by Barrett [1999] for finer-grained classifications and many references.

link), a (compound) system in this state has *no* particular value for the observable $F \otimes Q$ (where Q is the pointer-observable). The many-somethings interpretations go to the opposite extreme, and claim that *every* term on the right-hand side is ‘real’, corresponds to some reality. The apparatus has neither zero, nor one, of the values for the pointer-observable, but *all* of them.

Immediately, these interpretations face some questions. There is, of course, the obvious question of how systems can manage to ‘have’ multiple states — it appears to involve a straightforward contradiction. Generically, the reply is to introduce an indexical property that resolves the contradiction — the different values of the pointer-observable, for example, are realized *relative to* different values of the indexical property, which correspond to different *somethings* (‘worlds’, or in some versions, ‘minds’).

Everett’s original idea, however, seems rather to have been fundamentally relational in character. Indeed, consider a generic two-system state, as in (26). If the first system might be said, in some sense, to ‘really have’ the state $|x_j\rangle$, then *relative to* $|x_j\rangle$, the second system might be said to be in the state $|\psi_{\text{relative to } x_j}\rangle = K_j \sum_m c_{jm} |y_m\rangle$, where K_j is some constant of normalization. The probabilities generated by $|\psi_{\text{relative to } x_j}\rangle$ for the results of measurements on the second system are exactly those generated by the original compound state (for measurements that are restricted to the second system). Everett’s original idea seems to have been that systems possess states only ‘relative’ to the states of other systems. Others seem to adopt something like the view that the first system has various of the $|x_j\rangle$ in different ‘worlds’, and in those worlds, the second system has the corresponding relative state.

Each of these views raises difficult questions. In the case of Everett’s relational view, more needs doing. Consider the analogous case in space-time theories, where (prior to the 20th century), relationalists *asserted* that the only (spatio-temporal) reality is relational, but the only *theory* that was available (Newtonian classical mechanics) was *not* explicitly relational.⁷⁵ Everett’s relational view seems to be in somewhat of a similar situation. (For example, how and from what ‘point of view’ does the quantum state for the compound system itself get assigned? And how can *it* be understood relationally?) Moreover, the many-somethings views face the obvious metaphysical hurdle of making some sense of the plurality of ‘somethings’ (worlds, minds, whatever).

Here, I will raise just one additional question, commonly raised. The general prescription of these interpretations is to allow that *every* term in a superposition corresponds to a ‘reality’, in the sense just described. But recall that a given state can be decomposed in many ways, and that in general the terms in the superposition are different, depending on the decomposition. Which decomposition is ‘correct’?

One sort of answer to this question will postulate, or argue for, a *preferred basis*, in terms of which the decomposition is to be made. (The argument for a

⁷⁵See, for example, Brown and Pooley [2002] for details of the history, and see Butterfield, Ch. 1, §2.3, this volume, for further discussion of relationalism.

preferred basis is most often made on the grounds that decoherence (§5.4.5) picks it out, but there is a serious question whether the *near*-diagonality in position that decoherence typically effects is sufficient.⁷⁶) This answer is, in any case, highly problematic. For suppose that we had some good reason to believe that there is a preferred basis whose elements represented the true physical properties of all objects. In that case, we would hardly need the extravagance of the many-somethings interpretations to resolve the measurement problem. Relying on the preferred basis, we would simply assert that systems are always to be described in terms of this basis. There would then be no obstacle (for example, no Kochen-Specker contradiction) to defining probabilities over the properties picked out by this basis. (Indeed, below (§5.5.3.1) we will consider an interpretation — the de Broglie-Bohm theory — that can be characterized in precisely these terms.)

Another answer allows *all* decompositions to correspond to realities. Here, decoherence *can* be helpful, because it suggests that *creatures like us*, that is, creatures who interact with their environment in the particular way that we in fact do, will in fact be subject to decoherence, so that our perceptions will be correlated to the ‘right’ states of the objects of our perception, and we, as perceivers, will, as far as we can tell (with very high probability) be in worlds where the properties are the ones we typically believe objects to have. Of course, there remains to explain the point that in fact there are many ‘copies’ of each one of us (one for each value of the indexical), but we leave that issue to the advocates of these interpretations, who have their own ways of making sense of this type of plurality.

5.5.3 ‘Hidden-variables’ Interpretations

This basic idea behind any hidden-variables interpretation is that quantum theory (as EPR argued! see §4.4) is incomplete. These interpretations propose to supplement the state assigned to a system by quantum theory with an additional, ‘hidden’, state. (The variables need not in any sense be unavailable to observation; they are ‘hidden’ only from the eyes of quantum theory.) Of course, not just any proposal will do. In particular, such theories must somehow recover the empirical success of quantum theory. Doing so amounts to recovering the quantum probabilities as an average over the hidden states, in much the same way that one hopes to recover classical thermodynamics from classical statistical mechanics by averaging over the micro-states of the system. (See Uffink, Ch. 9, this volume.)

Hence, for example, label the hidden states by λ (where for simplicity we will assume that $\lambda \in \mathbb{R}$). Let $\rho_W()$ be the distribution of the hidden states in the (quantum) state W . Finally, let $\Pr_{\lambda,W}(F = f_k)$ be the probability assigned by the hidden state λ to the result (eigenvalue) f_k of a measurement of F on a system in the state W . (In the deterministic case, all such probabilities are of course 0 or

⁷⁶It is worth noticing as well that a state’s being *nearly* diagonal in position does not entail that the basis in which it *is* diagonal is anything close to position. See Bacciagaluppi [2000]. In these cases, which basis are we to say is ‘picked out’ by decoherence?

1.) Empirical adequacy then requires that

$$(104) \int_{\mathbb{R}} \rho_W(\lambda) \text{Pr}_{\lambda, W}(F = f_k) d\lambda = \text{Tr}[W P_{f_k}],$$

where P_{f_k} is the eigenspace of F corresponding to the eigenvalue f_k .

In this section, I will consider the de Broglie-Bohm theory (§5.5.3.1), a deterministic hidden-variables theory, and modal interpretations (§5.5.3.2), which are in general indeterministic.

5.5.3.1 The Theory of de Broglie and Bohm The de Broglie-Bohm theory is a deterministic theory of the trajectories of particles, somewhat in the mold of classical Newtonian (better, Hamiltonian) dynamics.⁷⁷ In its original form, the hidden-variables theory was formulated explicitly by separating the Schrödinger equation (including a potential, V , in the Hamiltonian) into a real and complex part, so that the solution, $\psi(x, t)$, is written as $\psi(x, t) = R(x, t)e^{iS(x, t)}$. The complex part has the form:

$$(105) \frac{\partial S}{\partial t} + \frac{(\nabla S)^2}{2m} + V - \frac{1}{2m} \frac{\nabla^2 R}{R} = 0.$$

This equation has the form of Hamilton’s equation, with the potential given by $V + U$ where U is the ‘quantum’ potential $-\frac{1}{2m} \frac{\nabla^2 R}{R}$. The real part has the form of a continuity equation, and implies the conservation of probability:

$$(106) \frac{\partial p}{\partial t} + \nabla \cdot p\dot{q} = 0,$$

where $p(x, t) = R^2(x, t) = |\psi(x, t)|^2$ is understood (stipulated to be) the probability of finding a particle at the point x (at time t) and

$$(107) \dot{q}(x, t) = \nabla S(x, t)/m$$

is interpreted as the velocity of the particle (when it is at the point x and time t).

Hence, one can understand the one-particle Schrödinger equation as describing an ensemble of particles (each of whose ‘quantum state’ is the same), distributed (in space) according to $p(x, t)$. By (106), this distribution is conserved over time, and thus may be understood properly as a probability distribution. The individual particles move as if they were classical particles governed essentially by Newton’s laws with the usual classical potential plus the additional ‘quantum potential’, U .

This theory is important for several reasons. Here I will mention three.

First, it shows that the eigenstate-eigenvalue link is far from a required principle of interpretation. Indeed, the eigenstate-eigenvalue link is quite strongly denied in

⁷⁷The basic idea was first suggested by de Broglie, most famously in his [1927]. Later (after de Broglie had been convinced that the theory does not work) it was developed by Bohm [1952], and has seen much further development in the past few decades. For a variety of approaches to and viewpoints on the theory, see Cushing, Fine, and Goldstein [1996].

this theory. The position observable does not even *have* eigenstates, and yet, on this theory, *every* particle *always* has a definite position.

Second, it provides a clear counter-example to oft-cited lessons of quantum theory, to the effect that quantum mechanics implies a lack of determinacy at the fundamental level, the impossibility of fundamental determinism, the impossibility of definite trajectories for particles, the impossibility of picturing quantum phenomena, and so on. The theory clearly describes particles with definite trajectories, moving deterministically under the potential $V + U$.

Third, this theory is probably the best-developed interpretation, with applications. There are compelling accounts of how particles behave during measurements of spin, for example (Dewdney, Holland, and Kyprianidis, [1986]). Moreover, the theory has occasionally been used to make headway on applications and problems that seemed much more difficult from the standard point of view. A nice example involves the prediction of how long a particle that will tunnel through a potential barrier will spend in the barrier. Determining this quantity is conceptually tricky (at best) in standard quantum theory, because the ‘particles’ have no definite trajectories, and there is no time observable in quantum theory. But the problem is conceptually straightforward within the de Broglie-Bohm theory (Leavens [1990]), precisely because one can essentially display the possible trajectories of a particle under the influence of the potential barrier (*and* the quantum potential), and then simply take the average of the times spent in the barrier in each trajectory.

However, we should not ignore some difficult issues faced by the theory. Setting aside (as outside the purview of this article) the (very significant) difficulty of extending the basic idea to relativistic quantum field theory, there are other questions that it faces. Here I will mention one.

Initial appearances notwithstanding, the ontology of the theory can stretch the imagination a bit. Indeed, one faces the following dilemma. On the one hand, if we take the ‘guiding field’ (quantum potential) to be a *real* potential field of some sort, it is highly non-local. (In the literature, this point is sometimes put in terms of the fact that it ‘lives’ not in \mathbb{R}^3 , but in the configuration space for the multi-particle system; of course, for a one-particle system they are the same.) This point gets reflected in the form of the ‘guidance condition’ (107) for many particles: $\dot{q}_i(x_1, x_2, \dots, x_n, t) = \nabla_i S(x_1, x_2, \dots, x_n, t)/m_i$. That is, the *velocity* of particle i depends not only on *its* location, but also on the locations of the *other* particles. Finally, notice that the particle does not respond to the *intensity* of the quantum potential in its region, but to the *form*. That is, increasing the intensity of the quantum potential (multiplying R in (105) by some constant > 1) does not change its effect on the particles at all. All of these (and more) observations will, for some, add up to the conclusion that the quantum potential is, at best, in need of some serious metaphysical clarification, and at worst, too bizarre to countenance.

On the other hand, one might then deny its reality. A particle moving according to the guidance condition, with $V = 0$ (zero ‘classical’ potential) will in general still ‘deviate’ from the *classically* expected trajectory, but are we thereby *required*

to suppose that this ‘deviation’ is due to the presence of some additional, non-classical, potential? Perhaps, instead, we can take the guidance condition as a fundamental rewriting of Newton’s laws.⁷⁸ In this case, we may avoid the oddness of the quantum potential as a physical entity, but the theory becomes extremely reductionistic — the *only* real property of particles is their spatio-temporal trajectory. Such a theory is forced to eschew providing explanations (for phenomena such as bonding between particles) that one might wish to have, and in general we *do* have in standard quantum theory.⁷⁹

5.5.3.2 ‘Modal’ Interpretations Modal interpretations⁸⁰ are in many ways similar to the de Broglie-Bohm theory: they also postulate hidden-variables (denying the eigenstate-eigenvalue link) and they also (at least, they *can*) formulate a dynamics for these hidden-variables. There are two major differences: (1) in general modal interpretations will allow that the physical quantities that are ‘definite’ for a system are state-dependent, and therefore they can change in time; (2) in general (and typically as a consequence of the choice they make for which physical quantities are definite) modal interpretations are stochastic. Many such interpretations have been proposed. Here we concentrate one class of such interpretations (‘spectral modal interpretations’) that has seen much development, and briefly refer to some more recent ideas.

Recall that every density operator can be uniquely decomposed as a weighted sum of its spectral projections. The central claim of spectral modal interpretations is that the set of observables with a definite value is the *largest* set that can consistently be assigned values, with the restriction that the spectral projections of the density operator must have values (0 or 1, in this case). In other words, treat the density operator as an observable, and assign it a value; then assign values to as much else as one can without running into a Kochen-Specker-type contradiction. It turns out that, under some reasonable assumptions, this idea leads, for any given state (density operator), to a unique set of observables that will get assigned a definite value.⁸¹ Because the state (of course) changes in time — and keep in mind that we are talking about the *reduced* density operator, so that it need not change unitarily — the definite-valued observables generally change in time.

Many (though not all) modal theorists aim to define some sort of dynamics for the definite properties of a system. The problem is complicated because there are

⁷⁸Indeed, the basic idea here is to reform the very notion of an ‘inertial’ trajectory. See Pitowsky [1991].

⁷⁹See Bedard [1999] and Dickson [2000].

⁸⁰There are many variants on the general theme of modal interpretations. The term itself was coined by van Fraassen [1972]. A renaissance of sorts occurred in the 1980s, with, for example, key works by Kochen [1985], Dieks [1988], and Healey [1989], and a later related but somewhat different view put forward by Bub [1997]. A helpful monograph is Vermaas [2000], and a helpful collection is Dieks and Vermaas [1998].

⁸¹As an example of this genre of theorem, see Clifton [1995]. Bub and Clifton [1996] prove a similar theorem, but geared towards Bub’s [1997] interpretation.

two sorts of dynamics occurring ‘together’: (deterministic) changes in the set of definite-valued properties (observables) for a system, and (stochastic) transitions from one definitely possessed property to another. It was shown by Bacciagaluppi and Dickson [1999] that the problem has a solution, and they exhibit some examples. However, just as in the de Broglie-Bohm theorem, there arises a serious difficulty about whether such a dynamics can be made Lorentz-invariant. Straightforward generalizations from the non-relativistic to the relativistic case are known to be unavailable, but it remains an open question whether a ‘natively’ relativistic version of the approach can work. The resolution of this issue is connected with the more general issue of whether, and if so how, the basic idea of the modal interpretation (‘make as many observables definite as possible, given certain constraints’) can be extended to quantum field theory. If it can, it seems that the most likely venue will be that of algebraic quantum field theory, for even the non-relativistic quantum-mechanical versions of the modal interpretation lend themselves nicely to an algebraic formulation.⁸²

Do modal theories solve the measurement problem? If so, they do because their prescription for choosing the definite-valued observables manages to pick out the observables — such as pointer observables — whose definiteness is threatened by the measurement problem. The de Broglie-Bohm theory (as well as the modal interpretation of Bub [1997]) accomplishes this feat more or less by fiat. So long as the properties we believe pointers (and cats, and so on) to have can ultimately be understood to supervene on the trajectories of the constituent particles, the de Broglie-Bohm theory has a convincing solution to the measurement problem. Modal interpretations must in general argue their way to this point. For example, is definiteness of the spectral projections of the reduced density operator enough to secure definiteness of the properties that we believe macroscopic objects have? This question, as it turns out, is quite difficult to answer. In the case of an idealized impulsive measurement, the reduced density operator for the apparatus *does* have the definite pointer-states as its spectral projections. But in realistic (non-ideal) measurements, there is potential trouble. At one time, it was thought that decoherence saves the day, but it is no longer clear whether decoherence does the work needed here.⁸³

5.5.4 Collapse Interpretations

5.5.4.1 Non-Dynamical Collapse Interpretation The idea that the quantum state must, in some sense or other, ‘collapse’, is quite natural as soon as one interprets it in terms of probability. And various suggestions have been made,

⁸²For example, see Halvorson and Clifton [1999]. For an initial foray into modal interpretations of quantum field theory, see Clifton [2000]. For a discussion of this foray especially as it relates to the issue of Lorentz-invariance, see Earman and Ruetsche [2006], which also includes references to the earlier work on Lorentz-invariance in modal interpretations. Cf. also Halvorson, Ch. 8, Section 5, this vol.

⁸³See Bacciagaluppi and Hemmo [1996] for the former attitude, and Bacciagaluppi [2000] for the latter.

from time to time, about when this collapse occurs. See, for example, Dirac [1930]

Consider an observation, consisting of the measurement of an observable α , to be made on a system in the state ψ . The state of the system after the observation must be an eigenstate of α , since the result of a measurement of α for this state must be a certainty.

That is, collapse occurs upon observation (i.e., measurement). The argument that Dirac is making here is the following. If we repeat the measurement of α on the system, we will get the same result that we got the first time (regardless of ψ , and assuming, of course, that the measurement was of the first kind — §3.1.2) with probability 1. Suppose that result was a . There is only one state that assigns probability 1 to a , namely, the eigenstate corresponding to a .

I have already suggested (§5.4.3) that an account along these lines is, at best, in need of help.⁸⁴ Otherwise, it is little more than a statement of the collapse postulate, with no clear account of what makes the difference between ‘measurements’ and other interactions.

One obstacle to drawing this distinction in a principled way is that if too *few* interactions count as measurements, then we might, after all, still be stuck with the measurement problem (because collapse might not occur in situations where we need it to get the definiteness of properties that we experience as definite). However, if too *many* interactions count as measurements, we might end up with a theory that is empirically false. (For example, if the state of a photon passing through the double-slit apparatus (figure 3) *always* collapses after it passes through the barrier, then we will never see an interference pattern, contrary to experimental results.)

One approach to drawing the distinction (between cases where collapse occurs and cases where no collapse occurs) that arguably avoids both of these pitfalls, and arguably does begin to smell like an interpretation of the theory, is due to Wigner [1961]. Wigner is motivated by the following argument, often referred to by the name ‘Wigner’s friend’.

Wigner asks us to imagine the following scenario. A measuring apparatus, μ , measures some observable, F (by means of a pointer-observable, M), on an object-system, σ . Meanwhile, both Wigner (α) and his friend (β) are in a position to observe μ (by means of ‘observation-observables’ A and B). Let us use $|m_n\rangle$ to refer to the eigenstates of M , and similarly for the rest.⁸⁵ Now, suppose that the measurement of σ by μ has occurred, while neither α nor β has observed μ . If σ was originally in the state $\sum_n c_n |f_n\rangle$ then we may write the state of the total system as

$$(108) \quad \sum_n c_n |f_n\rangle |m_n\rangle |a_0\rangle |b_0\rangle$$

⁸⁴And, alas, such help has only rarely been forthcoming. Hence I have no list of references to offer the reader, here. Perhaps the most famous attempt to make sense of non-dynamical collapse is precisely the one discussed here, due to Wigner.

⁸⁵There is a serious question whether a physical state such as ‘observes the pointer to be in the state $|M_n\rangle$ ’ is pure, and therefore whether there *are* observables such as A and B . See §1.4.3.

According to the eigenstate-eigenvalue link, in this state, μ and σ are *not* in a definite states of $M \otimes F$. At this stage, standard quantum theory will say to invoke the collapse postulate, so that the state becomes $|f_k\rangle|m_k\rangle|a_0\rangle|b_0\rangle$ for some k . In other words, in absence of specific knowledge about the outcome, the state becomes the (ignorance-interpretable) mixture

$$(109) \sum_n |c_n|^2 P_{f_n} \otimes P_{m_n} \otimes P_{a_0} \otimes P_{b_0}.$$

However, as noticed above (§5.4.5.2.a), there is a real empirical difference between (108) and (109), even if in general it is practically impossible to detect. Wigner finds nothing in the nature of σ , μ , or the interaction between them, to warrant the assumption that this physical change of state occurs.

Notice, moreover, that the observers α and β have not yet entered the picture in a substantial way. But it is *their* definite experiences that we are (in this argument) obligated to secure. Wigner considers himself (α), and notices that at *this* stage of the process (i.e., once the state is given by eq. 108), nothing threatens the definiteness of his experiences — in fact, he remains in the ‘has not yet observed’ state ($|a_0\rangle$) with probability 1. So consider the state after α observes the apparatus:

$$(110) \sum_n c_n |f_n\rangle|m_n\rangle|a_n\rangle|b_0\rangle.$$

Now Wigner feels compelled to assent to collapse, because otherwise (by the eigenstate-eigenvalue link) his own state would be indefinite.

But what about Wigner’s friend (β)? Suppose that β observes the apparatus before α . Wigner argues that, given a principle of charity about β ’s reports of prior mental states, he is obligated to collapse the state upon observation by β , even if he (α) has not yet observed μ . For suppose that α then (after β) observes μ , and then asks β : After you observed μ , did you feel that you experienced a definite outcome? Did you observe it to be in a definite state? Surely β will answer ‘yes’, and assuming (as Wigner does — this assumption is the principle of charity) that we are to believe β ’s report, we must assume that the state collapses upon observation by β .

Hence any observer to whom the principle of charity applies will effect a collapse of the state, upon observation. Wigner believes that the principle extends to anything that is, or has, a ‘mind’ (a term that is largely uncircumscribed, here). Moreover, given a fairly strong distinction between physical bodies and minds (i.e., given some form of dualism), one can point to the distinctness of minds as somehow the *reason* that observation by a mind collapses the state, while observation by a non-mind (e.g., a typical measuring apparatus) does not collapse the state.

Setting aside the question of the plausibility of the premises of Wigner’s argument, notice that we have arrived at this conclusion through the backdoor. In other words, Wigner does not give us an argument that proceeds from the nature of minds, and appeals to that nature to argue that observation by a mind will

collapse the physical state. Indeed, Wigner's view faces the same difficulty that all dualisms face: what is the connection between mind and matter? *How* does mind collapse the *physical* state of a system? Various extraordinarily speculative ideas have been floated, none particularly convincing. (See Atmanspacher [2004] for a review.)

5.5.4.2 Dynamical Reduction Theories Another strategy for making collapse well-defined is to describe it as a physical process. Of course, in one sense it already *has* been characterized as a physical process, for it is a change in the physical state of a system. Here, however, I mean something more, namely, formulating a single equation of motion that somehow incorporates *both* the continuous 'Schrödinger' evolution *and* collapse, in a unified way. There have been (and continue to be) many proposed schemes along these lines.⁸⁶ Here, again, I consider just one example.

5.5.4.2.a Intuitive Account of Continuous Spontaneous Localization

The example that I will consider is the continuous version of a theory originally proposed by Ghirardi, Rimini, and Weber [1986]. The continuous version, called 'continuous spontaneous localization' (CSL), is due to Pearle [1989]. It is easiest to state in terms of wavefunctions. The basic idea is that each physical system in the universe experiences, simultaneously, two types of evolution: the 'normal', deterministic, 'Schrödinger evolution', and a stochastic tendency towards localization. By the latter, we mean that the system's wavefunction (in configuration space) experiences random (but infinitesimally small) fluctuations, the net result of which is, on average and with overwhelming probability, to tend to make the wavefunction more localized — more of the probability gets concentrated in some localized region.

The strength of this stochastic evolution, which determines how quickly the localization occurs, is determined by the number of particles in the system.⁸⁷ For systems with microscopic numbers of particles, the stochastic part of the evolution is swamped by the Schrödinger evolution, so that the evolution is, for such systems, nearly identical to that given by standard quantum theory. But for systems with macroscopic (e.g., around 10^{23}) numbers of particles, if the system is initially in a superposition of two (or more) localized states whose centers of localization are far apart ($> 10^{-5}$ cm or so), then the stochastic element of the evolution acts very quickly to suppress (continuously) all but one of the terms in the superposition. The parameters of this stochastic process are arranged so that the probability of reduction to one state or the other is equal to the quantum probability assigned to that state (by the initial superposition).

⁸⁶See, for just a few examples, Diosi [1992], Ghirardi, Rimini, and Weber [1986], Gisin [1984], and Primas [1990].

⁸⁷There are schemes in which the strength of the stochastic evolution depends on other things, for example, mass.

5.5.4.2.b Mathematical Account of Continuous Spontaneous Localization Here are a few mathematical details. In CSL, the evolution of the wavefunction (here written in one spatial dimension) is given in terms of the (non-unitary)⁸⁸ operator:

$$(111) U(t) = \exp[-iHt] \exp \left[-\frac{\gamma}{2} t \int N^2(x) dx \right] \exp \left[\int N(x) B_x(t) dx \right].$$

The first term is just the usual quantum-mechanical evolution operator (for a time-independent Hamiltonian, H). The $N(x)$ are a family of ‘number density’ operators that, intuitively, indicate roughly the number of particles inside some region centered on x (or, in three dimensions, a ball of some fixed radius around x — this fixed radius is part of the definition of the $N(x)$, which we set aside here). The $B_x(t)$ is a field of stochastic processes (i.e., a continuous family of stochastic processes, each evolving at a point x). Hence the evolution operator in (111) is ‘stochastic’ — it depends on the values that the $B_x(t)$ take. Notice, though, that for a *given* realization of the $B_x(t)$, one can plug the realization into (111) and it would then be deterministic.

The second term in (111) is a ‘decay’ term, where the rate of decay depends on γ (and the particle-density). The third term in (111) is a ‘growth’ term, where the growth depends on the evolution of the stochastic process $B_x(t)$ (and the particle-density). These two terms ‘fight’ against one another, with the ultimate result tending (with extremely high probability) to be an overall decay of the wavefunction, except in some localized region, where it grows.

$B_x(t)$ is defined so that the likelihood of a growth in any given region is directly proportional to the amplitude of the wavefunction in that area. Hence the probability that the growth term will eventually cause the wavefunction to grow in a given area is indeed the same as the quantum-mechanical probability that a collapse would occur in that region. Notice that as the wavefunction begins to collapse onto some region, the probability that it will continue to collapse onto that region grows, because the amplitude of the wavefunction will be higher there, and therefore the stochastic processes in that region will be more likely to grow.

The final ‘trick’ is to choose γ so that the claim of the previous section is true, namely: for systems with a small number of particles, the first term in (111) dominates the other two (because $N(x)$ will be small everywhere), while for systems with a very large number of particles, the second and third terms will come into play. There are empirical constraints on the choice of γ , but it is in fact possible to find values for γ that are consistent with known experimental fact.

5.5.4.2.c Evaluation CSL has the obvious virtue of being a well-defined, mathematically precise, theory that manages to describe ‘collapse’ as a physical process in an unambiguous way. Nonetheless, several questions have been raised about it. Here I mention two.

⁸⁸In general, the norm of the wavefunction is not preserved by its evolution in CSL. However, it is a straightforward matter to renormalize the the wavefunction at any time.

The first concerns perhaps the biggest theoretical obstacle faced by the theory, namely, the formulation of a relativistic version. While progress has been made in various directions (e.g., Ghirardi, Grassi and Pearle [1990]), serious problems remain. Moreover, these problems are directly related to the characteristic feature of CSL, the introduction of the family stochastic processes $B_x(t)$, which, in the relativistic context, gives rise to infinities that do not appear in the standard theory (and cannot be ‘renormalized away’ in the same way).

The second issue concerns the fact that the collapse in CSL is never complete. That is, the state never becomes one whose support is *entirely* contained within a localized region. Instead, it always has ‘tails’ — the wavefunction always has non-zero amplitude at points outside the region in which the system has supposedly been localized.

The question here is whether one is still licensed to call such a system ‘localized’. If we maintain the eigenstate-eigenvalue link then the answer might well be ‘no’. After all, a wavefunction with tails is *not* an eigenstate of the position observable, nor any coarse-graining of it. The response on the part of at least some advocates of CSL has indeed been a new understanding of the wavefunction, as some sort of direct representation of, as Bell [1990] put it, ‘density of stuff’. This view, while avoiding the immediate problem, does have some metaphysically disturbing consequences, such as the fact that apparently there will, on this view, be very ‘faint’ (low-density) ‘copies’ of all objects located in many places.

6 NON-LOCALITY

Recall (§1.2.6.4) that compound systems in a non-factorizable state may exhibit correlations in the properties that they exhibit (upon measurement at least). These correlations may be ‘non-local’. Non-locality involves, first, the spatial separation of the particles (and the measurement-events on them that exhibit the correlated properties). There is the question, then, whether correlations exhibited by spatially separated systems can be explained in terms of a ‘common cause’, a single event in the history of the two systems that could explain the correlation between them.

In §6.1 I will clarify this question (§6.1.1), and then consider some theorems that show, under reasonable assumptions, that in fact no such common cause explanation is available (§§6.1.2–6.1.3). Hence quantum theory — and indeed the physical world itself! — appears to be ‘non-local’, in some sense or other. In §6.2 I will consider several reactions to these theorems, and the implications for the prospects of finding a relativistically invariant and also conceptually satisfactory interpretation of quantum theory.

6.1 No-Go Theorems

6.1.1 Non-local Correlations

6.1.1.1 Statistical Correlation The term ‘correlation’ comes from statistics, and it is being used here in that sense. Given two random variables, A and B , their correlation is defined to be:

$$(112) r_{AB} = \frac{E[(A - \bar{A})(B - \bar{B})]}{\sigma(A)\sigma(B)},$$

where $E[\cdot]$ is the expected value and $\sigma(\cdot)$ is the standard deviation. The correlation is 1 if $A = kB + m$ for positive k (and any fixed m , and -1 for negative k). The numerator of (112) is called the ‘covariance’. The denominator is present essentially for normalization.

It should be clear that correlation is a measure of the ‘dependence’ of the value of one variable on the other. Non-zero correlations will generally be a consequence of the fact that, for at least some possible values a and b of the random variables, $\Pr(A = a|B = b) \neq \Pr(A = a)$, i.e., A and B are not statistically independent.

6.1.1.2 Correlations in the Singlet State For a pair of spin-1/2 particles, the state

$$(113) |\psi_{\text{singlet}}\rangle = \frac{1}{\sqrt{2}} \left(|z_+\rangle|z_-\rangle - |z_-\rangle|z_+\rangle \right).$$

is called the ‘singlet’ state. (Notice that it is a two-particle analogue of the state in eq. 93.) Consider the observables S_z on each of the two particles. (One can think of these observables as random variables, because the state in (113) generates probabilities for the possible values of these observables, or more precisely, for the observables $S_z^{(1)} \otimes \mathbb{I}^{(2)}$ and $\mathbb{I}^{(1)} \otimes S_z^{(2)}$, where the superscripts number the particles.) The covariance of these observables (normalized to have eigenvalues ± 1 in order to bypass calculating standard deviations) in the state $|\psi\rangle$ above is:

$$\begin{aligned} r_{S_z^{(1)} S_z^{(2)}} &= \langle \psi | \left(S_z^{(1)} \otimes S_z^{(2)} \right) | \psi \rangle \\ (114) \qquad &= \frac{1}{2} \left(\langle z_+ | \langle z_- | - \langle z_- | \langle z_+ | \right) \left((-1) | z_+ \rangle | z_- \rangle - (-1) | z_- \rangle | z_+ \rangle \right) \\ &= -1. \end{aligned}$$

The values are perfectly anti-correlated. Indeed, a similar calculation for $S_u^{(1)} \otimes S_u^{(2)}$ (for an arbitrary direction u) shows that $r_{S_u^{(1)} S_u^{(2)}} = -1$; the perfect anti-correlation holds in every direction. For distinct directions, u and u' (but, without loss of generality given the spherical symmetry of the state, setting $\phi = \phi' = 0$ in (36) and again ‘normalizing’ the observables), we calculate:

$$(115) r_{S_u^{(1)} S_{u'}^{(2)}} = -\cos \theta \cos \theta' - \sin \theta \sin \theta' = -\cos(\theta - \theta').$$

This correlation is predicted by quantum theory, and well confirmed by a variety of experimental tests.

6.1.1.3 Common-Cause Accounts of Correlations Of course, in one sense, non-local correlations are ubiquitous. Consider: there is a correlation in the tides at different locations on the planet; there is a (negative) correlation between the average daily temperature in Sydney, Australia and London, England; there is a (positive, near-perfect) correlation in the number of words appearing in the newspaper at my front door each morning, and the one appearing at my neighbor's front door each morning. None of these correlations are surprising, even though they are correlations between the properties of spatially separated objects. The reason is clear: these 'non-local' correlations have a *local* explanation, in terms of a common cause.

Following Reichenbach [1956, 158-159], one can formalize the notion of a common cause in terms of conditional probabilities. Suppose that A and B are correlated. In this case, $\Pr(A \wedge B) \neq \Pr(A)\Pr(B)$. A common cause for the (let us say, simultaneously occurring) events A and B is an event, C , such that: (i) $\Pr(A|C) > \Pr(A|\neg C)$; (ii) $\Pr(B|C) > \Pr(B|\neg C)$; (iii) $\Pr(A \wedge B|C) = \Pr(A|C)\Pr(B|C)$; and (iv) $\Pr(A \wedge B|\neg C) = \Pr(A|\neg C)\Pr(B|\neg C)$. Conditions (i) and (ii) say that C is probabilistically relevant to the occurrence A and B respectively, while conditions (iii) and (iv) say that C 'screens A off from B ' (and vice versa) — that is, C completely accounts for the correlation between A and B .⁸⁹

Hence, for example, we may observe that the newspaper at my front door (A) was printed from the same master (C) as the one at my neighbor's front door (B). The same earth whose inclination towards the sun (C) partially determines temperature in Sydney (A) also partially determines temperature in London (B). And so on.

And yes, thanks to Einstein's theory of general relativity, even the moon's gravitational pull on the oceans is a *local* common cause: the gravitational 'force' propagates locally, as a wave, through space. Indeed, physics has, for at least the past century, been pushing towards *local* theories, and indeed theories that are, in some appropriate sense, Lorentz-invariant. (The precise meaning of 'Lorentz-invariant' varies, but the rough idea is that the theory does not permit the transmission of matter and energy faster than light, or does not permit a signal to be sent faster than light, or at the very least, does not permit matter to be accelerated from subluminal to superluminal speeds. See Malament, Ch. 3, §2, this volume.)

Quantum theory permits non-local correlations. The question we face is whether those correlations can also, like the others mentioned here, be given a (local) common cause explanation. In particular, can the correlations implied by (115) be explained by a common cause?

⁸⁹This formulation is not entirely satisfactory, but illustrates the notion of a common cause well enough for present purposes. See Uffink [1999b] and references therein.

6.1.2 Bell's Theorem

Bell's theorem essentially answers 'no' to this question. Moreover, Bell was able to derive some constraints on the predictions made by any local theory. These constraints are quite convincingly violated by experiment.

6.1.2.1 The Experiment The experiment considered by Bell is essentially the one considered by EPR (§4.4), in a version described by Bohm [1951, ch. 22]. There is a source of pairs of spin-1/2 particles, each pair in the state $|\psi_{\text{singlet}}\rangle$ (113); the particles are directed towards Stern-Gerlach devices, each oriented in some direction (labeled u and u'). The particles arrive at these devices in such a way that the measurements are made at space-like separation. Indeed, even the choices of the directions u and u' are made at space-like separation. That is, the choice of measurement, and the measurements themselves, are made sufficiently far apart in space, and sufficiently close together in time, that a signal of any sort from one measurement-event to the other would have to travel faster than light. After many runs of the experiment, the results are collected, and correlations determined.

In fact, in order to get Bell's Theorem underway, we need consider only three possible directions (u_a , u_b , and u_c) on each side. The measurement-devices thus choose, at the 'last second', which of these three directions to measure, and then they record the result.

6.1.2.2 The Issue in Terms of 'Hidden Variables' One way to think about the question about common causes that we raised above is in terms of so-called 'hidden-variables theories' (§5.5.3) and indeed Bell conceived of the issue in these terms. Recall that hidden variables describe complete states of the particles, states that include information not provided by the quantum state. In the deterministic case, these complete states fix the values of S_u for each value of u , for each particle. The picture, then, is that, back at the source, the pair of particles about to be released towards the two measurement-devices has some 'hidden' state. This state already determines, for each particle, and each possible direction u_a , u_b , and u_c , a value (+1 or -1) for spin.

Indeed, recall the EPR argument (§4.4), and their conclusion that quantum mechanics is 'incomplete'. The point there was that locality (plus the criterion for physical reality) forces one to conclude, in the context of this experiment, that each system must in some way *already have* values for all S_u — they must have 'established' the correct correlations back at the source (for they cannot do so later, at spacelike separation, by locality).

Note that for a given direction, the hidden state must determine the particles to have *opposite* spin, because of the perfect anti-correlations. (These hidden states can thus be partitioned into eight classes, two possible assignments of spin for each of the three directions.) Apart from that, the only requirement is that it be possible to recover the quantum-theoretic correlations by averaging over some

distribution of the hidden states, as in (104).

6.1.2.3 Probabilistic Hidden States A slightly more general approach is to allow the hidden states to determine spins only probabilistically. The theory will remain ‘local’ as long as the hidden state of the particles back at the source screens off the measurement-event at one side from the results at the other side (§6.1.1.3). Letting Pr_λ be the probabilities for measurement-results prescribed by the hidden state λ (and we are assuming, now, that the quantum state is fixed as $|\psi_{\text{singlet}}\rangle$), the screening-off condition is that, for any $u, u', k = \pm 1$, and $k' = \pm 1$:

$$(116) \text{Pr}_\lambda(x = k|i = u, j = u', y = k') = \text{Pr}_\lambda(x = k|i = u),$$

where x is the result for particle 1, y is the result for particle 2, i is the direction of spin measured on particle 1, and j is the direction of spin measured on particle 2. The intuitive idea, here, is that the result for particle 1 depends only on the direction of spin measured (and λ), and is thus independent of the result for particle 2, and the direction in which spin was measured on particle 2. This condition, or something similar, is often called ‘Bell locality’.

So, finally, the question we face is whether there can be hidden states, λ , and a distribution over them, $\rho()$, such that the probabilities, Pr_λ both obey Bell locality (116) and reproduce the quantum correlations (115) via averaging as in (104). (Note that a similar condition will hold for the single-wing probabilities, i.e., those obtained for just *one* of the two particles, which will also, of course, be obtained by averaging over hidden states.)

In general, as I mentioned earlier (§5.5.3), the distribution $\rho()$ would depend on the quantum state (which, recall, we are presuming here to be $|\psi_{\text{singlet}}\rangle$). But it must not depend on anything that would require a non-local influence on the source, because $\rho()$ is supposed to represent the distribution of hidden states amongst pairs produced at the source. For example, we will presume that the choice of measurement-direction and the results of the measurements do not influence $\rho()$.

6.1.2.4 Bell’s Theorem Bell’s theorem shows that in fact no common-cause explanation of the sort outlined above exists.

Theorem (Bell [1964]): For any choice of directions u_a, u_b , and u_c , any hidden-variables theory that satisfies (i) Bell locality (116) and (ii) independence of $\rho()$ from the measurement-events will obey the inequality

$$(117) |r_{ac} - r_{bc}| - r_{ab} \leq 1.$$

(where $r_{ab} := r_{S_{u_a}^{(1)} S_{u_b}^{(2)}}$ and so on).⁹⁰

⁹⁰Numerous other inequalities have since been derived, with various virtues, especially regarding experimental testability. See Clauser and Shimony [1978].

Given the quantum-mechanical prediction (115), this inequality is violated by many choices for u_a , u_b , and u_c . For example, defining each u_k in terms of Euler angles ϕ_k and θ_k , choose $\phi_k = 0$ for $k = a, b, c$ (so that these directions are coplanar), $\theta_a = 0$, $\theta_b = \pi/4$, and $\theta_c = \pi/2$. In this case, the quantity on the left-hand side of (117) is $\sqrt{2}$.

6.1.2.5 Experimental Violation of Bell's Inequality Bell's Inequality is violated by experiment. Of course, as is always the case with experimental realizations of theory, there are complications. For example, actual sources often do not produce particles that move away from the source in precisely opposite directions. One must either filter out those that do not, or take account of the fact that they do not move in exactly opposite directions. Neither solution is completely straightforward. And there are other issues with existing experimental tests of Bell's inequality, although experimentalists are increasingly able to settle these issues in favor of the violation of the inequality in the world, i.e., the confirmation of the quantum correlations.

There is one issue that will not be settled by experiment, at least not entirely, and it involves the fact that for any given pair of particles, only *one* pair of measurements (one on each particle) can be made. In other words, for each given pair, we are not *really* probing the entire hidden state, but only its implications for one pair of measurements. We must therefore *presume* that the sample we are getting of ensemble of hidden states is representative. To put the point another way, consider all of the pairs where we measure, for example, u_a, u_b , and then all the pairs where we measure u_b, u_c . In order to test Bell's inequality, we will calculate the correlations amongst these two groups of pairs, then plug the results into (117). In doing so, we are assuming that *had* we measured u_b, u_c on the u_a, u_b -pairs (i.e., *instead* of measuring u_a, u_b on them), we *would* have gotten the same correlation that we in fact got for our actual u_b, u_c -measurements. Most consider this counterfactual innocuous — in large part because violating it seems to involve one in some rather odd conspiracy theories, many of which would in any case be non-local.⁹¹ It is worth pointing out, however, that the counterfactual nature of the assumption is not due to contingent facts about what we happened to measure (as the corresponding counterfactuals often would be in classical physics), but rather it is due to the in principle incompatibility of the measurements involved — S_{u_a} and S_{u_b} *cannot* be measured at the same time on the same particle, and, moreover, in accordance with the uncertainty principle (§4.2), measuring one destroys whatever previous knowledge we have of the other.

⁹¹For example, it will not do to suppose that only a 'certain kind' of hidden state is ever involved in a u_b, u_c -measurement, because this theory would violate the condition that the distribution at the source be independent of the choice of measurement. However, models that rely on detector inefficiencies can assert — however oddly — that detectors have a preference for revealing only 'certain kinds' of hidden states, based on which measurement was performed. Fine's 'prism models' are of this sort — see Fine [1991] and Szabó and Fine [2002]. Note that detector efficiencies continue to improve, and could eventually rule out such theories on purely empirical grounds.

Finally, note that although Bell's Theorem is framed in terms of so-called 'hidden' states, it is ultimately *not* a theorem about hidden-variables theories, but about *locality*. After all, there is nothing stopping us from taking the λ to be the quantum state itself and following through Bell's derivation. Of course, in doing so, we will make an assumption (Bell locality) that is violated by quantum theory, and we could have checked this violation directly if we had wanted to. Either way the conclusion is the same: in the sense discussed here (Bell locality), standard quantum theory is itself a non-local theory.

6.1.3 Other Bell-Like Theorems

Plenty of other theorems are floating around that also bear on the issue of locality. A couple of them we have already seen (§5.4.1.2 and §5.4.1.2.c), and so will be mentioned only briefly here.

6.1.3.1 The Kochen-Specker Theorem as a Theorem about Locality

The non-contextuality condition in the Kochen-Specker theorem can, in the right circumstances, be seen as a locality condition, and the sought-for homomorphisms (recall §5.4.1.2.b) from the lattice of subspaces to the Boolean lattice $\{0, 1\}$ are the 'hidden states'. However, rather than develop this point in the context of the general Kochen-Specker theorem, let us just consider the point in the context of the GHZ theorem, which in fact was originally (and correctly) characterized by its authors as a Bell-like theorem without inequalities — one need only suppose that the three particles are space-like separated.

In this context, the condition of non-contextuality becomes this: the value assigned to each of the observables in (99) must be assigned independently of which observables (100) one considers it to be a part, that is, independently of the 'context of measurement', which is given by the set of observables measured on each of the three particles. In other words, the value that particle 1, for example, has for $S_x^{(1)}$ cannot depend on whether we are measuring $S_x^{(1)} \otimes S_y^{(2)} \otimes S_y^{(3)}$ or $S_x^{(1)} \otimes S_x^{(2)} \otimes S_x^{(3)}$. This condition is a consequence of locality, because such a dependence would imply that the result of the measurement on particle 1 would depend on which observable we decided to measure on particles 2 and 3, even if those decisions were made at space-like separation from the measurement-event on particle 1.

6.1.3.2 Hardy's Argument

GHZ simplified the conclusion of Bell's argument, getting rid of the inequality, but at the expense of introducing another particle. Hardy [1992] managed to make an argument not based on inequalities, but considering only two particles.⁹² He considers the experimental arrangement depicted in figure 5, which shows a 'double-interferometer' (called 'double' because there are two interferometers, 'overlapping' at the point A). The electron (e^-) and

⁹²The exposition here follows that in Dickson [1998, pp. 209–211].

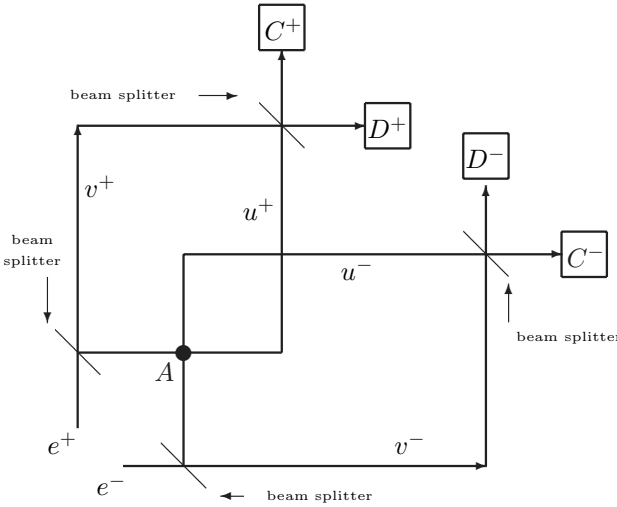


Figure 5. Hardy's Experiment.

positron (e^+) enter the interferometers on the paths s^- and s^+ respectively. They encounter a beam-splitter, and as a result, follow either the 'u' path or the 'v' path, with equal probability. If they both follow their u path then they meet at A and annihilate one another. Otherwise, they both reach a second beam splitter and the positron then goes on to either the detector C^+ or the detector D^+ , with equal probability. The electron is in the corresponding situation.

Hardy assumes that a hidden-variable theory will assign a definite path (u or v) to each particle. He claims further that the question whether a particle is on a given path ought to have a Lorentz-invariant answer — its path ought not depend on an observer's reference frame. Otherwise the hidden-variable theory is not Lorentz-invariant, and is, in that sense at least, non-local. Here, we will consider three reference frames. In the first, F^+ , the detection of the positron (at C^+ or D^+) occurs before the electron has passed through its beam splitter. In the second, F^- , the order of these events is reversed. And in the third, the detections are simultaneous.

The initial state of the electron and positron is $|s^-\rangle|s^+\rangle$ respectively, indicating that (before they reach the initial beam-splitter) they are on the paths s^- and s^+ respectively. The initial beam splitters effect the evolution

$$(118) |s^\pm\rangle \rightarrow (1/\sqrt{2})(i|u^\pm\rangle + |v^\pm\rangle).$$

Similarly, the second set of beam splitters effect the evolution

$$(119) |u^\pm\rangle \rightarrow (1/\sqrt{2})(|c^\pm\rangle + i|d^\pm\rangle)$$

$$(120) |v^\pm\rangle \rightarrow (1/\sqrt{2})(i|c^\pm\rangle + |d^\pm\rangle).$$

Using (118), the state after the system passes the point A is (in frame F)

$$(121) \frac{1}{2} \left(-|\gamma\rangle + i|u^+\rangle|v^-\rangle + i|v^+\rangle|u^-\rangle + |v^+\rangle|v^-\rangle \right),$$

where $|\gamma\rangle$ is the state after an annihilation. One can calculate the state at various times in the frames F , F^+ , and F^- by applying one or more of (118–120). I leave these calculations to the reader.

Now, in F^+ , if the positron is detected at D^+ , then the electron must be on the path u^- (with probability 1). Similarly, in F^- , if the electron is detected at D^- , then the positron is on the path u^+ . However, in F , the state of the pair before either particle goes through its second beam splitter is orthogonal to a state where the particles take the paths u^+ and u^- (because of the annihilation in that case).

Consider, then, a run of the experiment in which detectors D^+ and D^- both register a hit. (This outcome does indeed have non-zero probability.) In that case, an observer in F^+ will conclude that the electron took the path u_- . An observer in F^- will conclude that the positron took the path u^+ . *Both* of these statements must be true for an observer in F , but as we saw above, they cannot both be true. Hence, Hardy concludes, a theory that assigns definite paths to the particles cannot be Lorentz-invariant.

Note, finally, that Hardy has implicitly assumed non-contextuality as well. A hidden-variables theory is obliged only to reproduce the *experimental* predictions of quantum theory. If the theory is contextual, then in particular the path of a particle may depend on the presence (or not) of a device on the path that will detect the particle (or not) on the path. Hence, for example, such a theory is *not* obliged to assign the path u^+ to the positron if the electron is detected at D^- (in F^-) *unless* there is a detector in place to determine whether the positron is on u^+ . But in this case, the detector will interact with the positron, and everything changes, including (most importantly) the quantum-theoretic calculations that are behind Hardy's argument. In the absence of such a detector, and such an interaction, a hidden-variables theory need not respect the quantum probabilities. But agreed, a non-contextual theory *does* always have to respect those probabilities, because such a theory cannot alter the (hidden) state (i.e., the path) that it assigns to a particle depending on whether a detector is in place along the path u^+ .

6.2 Reactions to the Theorems

What should we make of the failure of locality? In this section, I shall consider four more precise versions of that question.⁹³ Can the failure of locality be used to send signals from one measurement-station to the other (§6.2.2)? Does the failure of locality imply the existence of some causal connection between the measurement-stations (§6.2.3)? Does the failure of locality imply (or can it be understood in terms of the claim that) the two systems are somehow not really distinct (§6.2.4)?

⁹³Extensive discussion of these issues and more is available in many places. Cushing and McMullin [1989], Butterfield [1992], Maudlin [1994], and Dickson [1998, chs. 6-9] will get the interested reader started.

Finally, and perhaps most important from the point of view of theoretical physics, what does the failure of locality imply about the possibilities for a fully Lorentz-invariant (relativistic) quantum theory (§6.2.5)?

Prior to asking these questions, I shall review (§6.2.1) an important analysis of locality, one that some have claimed helps resolve some of these questions.

6.2.1 *Parameter Independence and Outcome Independence*

Once we have accepted the non-locality of quantum theory, and of empirically adequate hidden-variables theories — and it is surprising how far some will go to resist this conclusion — the question becomes how to understand Bell locality. The most famous analysis of this condition is in terms of two other conditions, often called ‘parameter independence’ and ‘outcome independence’. The former expresses the idea that the parameters — the settings on (direction of spin to be measured by) the measurement-apparatus — at one measurement-station do not affect the outcomes at the other measurement-station. The latter expresses the idea that the outcomes at one measurement-station do not affect the outcomes at the other.

These conditions are, like Bell locality, statements of probabilistic independence:⁹⁴

Parameter Independence for all i, j, k, k', λ ,

$$(122) \Pr_{\lambda}(x = k|y = k', i = u, j = u') = \Pr_{\lambda}(x = k|y = k', i = u)$$

and similarly, reversing the roles of particles 1 and 2.

Outcome Independence for all i, j, k, k', λ ,

$$(123) \Pr_{\lambda}(x = k, y = k'|i = u, j = u') \\ = \Pr_{\lambda}(x = k|i = u, j = u') \times \Pr_{\lambda}(y = k'|i = u, j = u').$$

(See 6.1.2.3 for the notation.) The conjunction of Parameter Independence and Outcome Independence yields Bell locality with just some trivial applications of probability theory.

One of the values of this analysis is that it reveals, more precisely, how quantum theory violates Bell locality. In particular, quantum theory violates Outcome Independence, but satisfies Parameter Independence. The derivation of the correlation (115) more or less proves the first claim.⁹⁵ We will see one way of proving the second claim in the next subsection.

⁹⁴There are subtly different versions of these conditions, taking into account, for example, hidden variables in the apparatus, and so on. We are skating over such distinctions here. In addition, there are hosts of other locality conditions one might introduce. See Dickson [1998, chs. 6-9]. Jarrett [1984] is responsible for the first clear statement of this distinction. The conditions as stated here, and the terminology used here, is due to Shimony [1986].

⁹⁵Strictly speaking, in order to implement the probabilities in (123), one should introduce the

6.2.2 Signaling

6.2.2.1 Locality and Signaling Part of the original intention behind the analysis of Bell locality into Parameter Independence and Outcome Independence was to isolate a part of Bell Locality that does *not* involve a violation of relativity theory, in some sense or other. In particular, the claim is often made that a failure of Outcome Independence is somehow consistent with relativity, while a failure of Parameter Independence is not.

In particular, one might understand relativity theory as prohibiting superluminal signaling. A violation of Outcome Independence, so the argument goes, does not involve the possibility signaling, because even though it implies that the outcome at one measurement-station depends probabilistically on the outcome at the other (and of course the measurement-events are space-like separated), the outcomes themselves are probabilistic. That is, an experimenter cannot control the outcomes. But controlling the outcomes would be necessary in order to use the dependence between outcomes to send a signal.

On the other hand, experimenters *are* in control of parameters — they are in fact normally assumed to be the result of a free choice of the experimenter. Hence, so the argument goes, a violation of Parameter Independence implies that an experimenter can (probabilistically) influence the outcomes at the other station by manipulating the parameter (direction in which spin is measured) at the local station.

However, keep in mind that the probabilities in Parameter Independence and Outcome Independence are those generated by the hidden state, λ . If the experimenter is not in control of these hidden states, then a failure of Parameter Independence will *also* not imply the possibility of signaling. Moreover, control of the hidden states would mean that in fact a violation of Outcome Independence *also* implies the possibility of signaling, so long as the probabilities for the outcomes generated by different hidden states are different. In this case, an experimenter *can* in fact locally manipulate the outcomes in the probabilistic sense of controlling the hidden state in order to make a given outcome more or less likely, and thus, via the probabilistic dependence between outcomes, influence (probabilistically) the outcomes at the other station. outcome independence

An illuminating example of the significance of being able to control the hidden state is provided by the de Broglie-Bohm theory. It is clear that the theory satisfies Outcome Independence, because it is a fully deterministic theory. That is, given the parameters i and j , and the initial state (which, in this theory, is the initial positions of the particles, plus the quantum wavefunction), the results of the measurements are fixed. Hence, in particular, the result of one measurement, being fixed by λ , i , and j , does not depend on the result of the other measure-

apparatuses as two additional systems, and define the apparatus-settings (choice of direction of spin) as states of these additional systems. Hence one would have a state in a four-fold tensor product Hilbert space, and would consider the probabilities generated by that state. Following through the analysis in this way, however, is not particularly enlightening, and yields the same result, namely, the failure of Outcome Independence.

ment. On the other hand, the de Broglie-Bohm theory *does* violate Parameter Independence, for the result of a measurement on one particle *does* in general depend on the direction in which spin is measured on the other. (Changing the setting on the apparatus changes the wavefunction for the compound system in a way that ultimately changes the quantum potential, and hence the trajectories for *both* particles.) However, knowing *how* a particle is affected by a change in the parameters (and therefore being able to control this effect) requires knowing its precise position but recall that the precise positions of the particles are *not* known in this theory. Hence the violation of Parameter Independence cannot be used for signaling.

Indeed, as Cushing [1994] points out in this context, Valentini [1991a; 1991b] has shown that in the de Broglie-Bohm theory, signaling is possible if and only if the distribution of the particles is *different* from that given by the usual quantum-mechanical probability. Hence, given that Bohm's theory *does* respect this distribution, we have a case where, phenomenologically, the theory violates Outcome Independence and respects Parameter Independence, while at the level of the (uncontrollable!) hidden-variables (trajectories of particles), the reverse is true.⁹⁶ Hence one should be very cautious about identifying a failure of Parameter Independence with the possibility of signaling.

6.2.2.2 The No-Signaling Theorem While the distinction between Outcome Independence and Parameter Independence is, as explained above, not the same as the distinction between the impossibility and possibility of signaling, it is true that quantum theory (which, recall, obeys Outcome Independence but not Parameter Independence) does not permit signaling from one measurement-station to the other (or, indeed, non-local signaling in general). This claim is the conclusion of the quantum no-signaling theorem.

There are (at least) two routes to this conclusion, one focused on the effect that local generic interactions can have on the states of spatially separated systems, and one focused on the effect that measurements can have on the results of measurements on spatially separated systems. I shall outline each in turn. In both cases, we will consider a two-particle system in the (generally entangled) state W , and suppose that the two particles are spatially separated.

6.2.2.2.a Local Interactions and Reduced States Now consider the evolution of the compound system by two different unitary evolutions (generated by two different Hamiltonians), $U^{(1)} \otimes U^{(2)}$ and $U^{(1)} \otimes \tilde{U}^{(2)}$. The no-signaling theorem says that the reduced state for particle 1 is the same at the end of the evolution, whether the system evolves according to $U^{(1)} \otimes U^{(2)}$ or $U^{(1)} \otimes \tilde{U}^{(2)}$. In

⁹⁶In fact, Valentini's [1991a; 1991b] view, supported by a kind of quantum analogue of the classical H-Theorem (see Uffink, ch. 9, §I.4.2, this volume), is that the standard quantum distribution is a kind of 'equilibrium' distribution, to which the universe naturally relaxes. It is possible (indeed quite likely), on this view, that the universe is not *quite* in the equilibrium distribution, which means that in principle one could signal, if one could find an ensemble of particles not in the equilibrium distribution.

fact, under the first evolution, the reduced state for particle 1 is (recall §1.2.6.3.a)

$$(124) \quad W^{(1)} = \text{tr}^{(2)}[U^{(1)} \otimes U^{(2)} W(U^{(1)} \otimes U^{(2)})^{-1}] = \sum_n \langle e_n | U^{(1)} \otimes U^{(2)} W(U^{(1)} \otimes U^{(2)})^{-1} | e_n \rangle.$$

where $\{|e_n\rangle\}$ is some orthonormal basis in the Hilbert space for particle 2. But recall that the (partial) trace functional does not depend on the choice of this basis. Hence, when figuring the reduced state under the alternative evolution, $\text{tr}^{(2)}[U^{(1)} \otimes \tilde{U}^{(2)} W(U^{(1)} \otimes \tilde{U}^{(2)})^{-1}]$, we need only choose a basis, $\{|\tilde{e}_n\rangle\}$, such that $\tilde{U}^{(2)}|\tilde{e}_n\rangle = U^{(2)}|e_n\rangle$. Because $U^{(2)}$ and $\tilde{U}^{(2)}$ are unitary, such a basis always exists.

In other words, no interaction with just one of the systems can effect the reduced state of the other, and because the marginal probabilities for particle 1 (the probabilities for outcomes of measurements of observables on particle 1 only) depend only on the reduced state $W^{(1)}$, no interaction that involves just particle 2 can change the statistics of measurement-results for particle 1 (and, of course, vice versa).

6.2.2.2.b Single-System Measurements But what about measurements?

Consider a measurement of an observable $\mathbb{I} \otimes G$, i.e., a measurement of G on system 2, leaving system 1 untouched. Let $G = \sum_n g_n P_n$. Of course, we have already seen that if we project the state of the compound system onto the result of this measurement, then in general (for example, if W is the singlet state), the probabilities for results of measurements on system 1 will change. That fact is just a rehearsal of the existence of correlations between the two systems. But keep in mind that we are talking about the possibility of *signaling*, here. An observer who is restricted to the vicinity of system 1 will not know the results of measurements on system 2, but will know only that a measurement of $\mathbb{I} \otimes G$ might or might not occur. To such an observer, the measurement of $\mathbb{I} \otimes G$ on particle 2 is a ‘non-selective’ measurement, meaning that the most that this observer can say about the state after the measurement is that the compound system is now in the state

$$(125) \quad \sum_n (\mathbb{I} \otimes P_n) W (\mathbb{I} \otimes P_n).$$

It is instructive to see why (125) is the state after a non-selective measurement. Suppose that the result were in fact g_k . In that case, applying the collapse postulate, the state would be $P_k W P_k / \text{Tr}[W P_k]$. (Compare this expression with Lüder’s Rule, eq. 19.) Now consider that the probability of the result P_k is $\text{Tr}[W P_k]$, so that if we do not know the result of the measurement (i.e., the measurement is non-selective) then the state is a weighted sum of all possible results, the weights given by the probabilities of the various results. I.e., the state is (125).

The idea here is that an experimenter (‘B’) in the vicinity of particle 2 will attempt to send a signal to an experimenter (‘A’) in the vicinity of particle 1 by choosing to measure G (on particle 2) or not. The question, then, is whether A

can detect any change in the statistics of results for measurements on particle 1 as a result of B 's measurement on particle 2.

The answer is 'no'. To see why, consider an observable, $F \otimes \mathbb{I}$, and let $F = \sum_n f_n Q_n$. In the state W , the probability for the result f_m (of a measurement of $F \otimes \mathbb{I}$) is

$$(126) \text{Tr}[(Q_m \otimes \mathbb{I})W].$$

Suppose, instead, that the state is (125), i.e., a non-selective measurement of $\mathbb{I} \otimes G$ has been made. In that case, the probability for the result f_m (of a measurement of $F \otimes \mathbb{I}$) is

$$(127) \text{Tr} \left[(Q_m \otimes \mathbb{I}) \sum_n (\mathbb{I} \otimes P_n) W (\mathbb{I} \otimes P_n) \right].$$

By the linearity of $Q_m \otimes \mathbb{I}$, it can be taken inside the sum, and by the linearity of the trace functional, the trace of a sum then becomes a sum of traces. Moreover, because $Q_m \otimes \mathbb{I}$ commutes with $\mathbb{I} \otimes P_n$ (for any n, m), this sum becomes:

$$(128) \sum_n \text{Tr} [(\mathbb{I} \otimes P_n)(Q_m \otimes \mathbb{I})W(\mathbb{I} \otimes P_n)].$$

Using (12) together with the fact that $PP = P$ for any projection P , (128) is

$$(129) \sum_n \text{Tr} [(\mathbb{I} \otimes P_n)(Q_m \otimes \mathbb{I})W].$$

Taking the sum back inside the trace functional and observing that the $\mathbb{I} \otimes P_n$ constitute a resolution of the identity (i.e., $\sum_n \mathbb{I} \otimes P_n = \mathbb{I} \otimes \mathbb{I}$), we find that (129) is exactly (126). In other words, no measurement that A can make on particle 1 can determine whether the compound state is the original W , or the state (125) — B 's non-selective measurement of $\mathbb{I} \otimes G$ has no effect on the probabilities for outcomes of measurements performed by A on particle 1. Hence B cannot send a signal to A by means of such a measurement.

6.2.3 Causation

Does a failure of Bell locality (or one of its constituent conditions) imply a causal connection between the two measurement-stations? This question is much-discussed and disputed, and I certainly will not settle it here. But let us consider a few approaches.

Consider a counterfactual approach to the analysis of causation, one that, for example, takes as a sufficient condition for causal connection the truth of counterfactuals such as 'if it had been that C then it would have been that E ' and 'if it had not been that C , then it would not have been that E '.⁹⁷ It would appear, adopting some such understanding of causation, that causal connections *do*

⁹⁷Causation has frequently been associated with counterfactuals such as these. Lewis [1973] is famous for such an analysis, and as his work clearly shows, one must add several additional conditions to this basic idea — for example, conditions about the similarity of possible worlds, where those conditions will effect the semantics, hence truth-values, of the counterfactuals involved.

exist between outcomes, or between parameters and outcomes, in a deterministic hidden-variables theory that violates Outcome Independence, or Parameter Independence, respectively. If the hidden-variables theory is probabilistic, then perhaps one would conclude instead that there is a relation of probabilistic causality, though of course one would have to formulate appropriate probabilistic counterfactual conditions, for example, replacing ‘it would [not] have been the case that E ’ with ‘the probability of E would have been higher [lower]’.

On the other hand, others prefer an account of causation that requires that causal connections are underwritten by ‘causal processes’ that can transmit a ‘mark’.⁹⁸ That is, impressing some sort of ‘mark’ on the cause should leave a trace ‘from the cause to the effect’, and ultimately in the effect itself. There are different understandings of what this account really means, but on at least some understandings — most clearly, those that require the mark to travel continuously in space — the violation of Parameter Independence and Outcome Independence does *not* imply the existence of a causal connection between the measurement-stations.

Finally, some, e.g., Collier [1999], have argued that causation amounts to the transfer of information, in the information-theoretic sense. Maudlin [1994, ch. 6] has argued that there *is* a transfer of information in the Bell-type experiments. Hence, if one buys these arguments, there *is* causation between the measurement-stations, on this account of causation.

6.2.4 Holism

The issues surrounding Bell’s Theorem have led some to a radical view of entangled states as representing ‘holistic’ properties of entangled pairs (or n -tuples) of particles. On the one hand, we can agree that the properties represented by, for example, the (projection onto the) singlet state are not reducible to, and do not supervene on, properties of the individual particles. (Otherwise, a local hidden variables theory *would*, after all, be possible.) On the other hand, some claim that these holistic properties *explain* our inability to provide a common cause explanation for the quantum correlations in a way that does not violate any version of locality that we should care about, and does not, in particular, imply any violation of relativity theory. The basic idea is clear enough: the ‘two’ (or more) particles involved in a Bell-type experiment are not *really* ‘two’ particles; they are really one object, and therefore, no matter how far apart they may be, there is no question of an ‘influence from one to the other’, for there is no ‘one and other’ but only the whole thing.

There are, however, some *prima facie* problems with this proposal. First, it is not clear that it really explains anything; perhaps it only restates the problem in the form of an equally mysterious doctrine. Moreover, it is unclear what the real difference is, as regards the apparent conflict with relativity theory, between a space-like influence between two distinct objects and a space-like ‘connection’

⁹⁸See, for example, Salmon [1984].

between different ‘parts’ of the same object. (Of course, we must be very cautious about the meaning of ‘part’ here.) Indeed, the events that are involved in Bell-type experiments are well-defined, localized, events in space-time, viz., macroscopic pointer-readings. There are correlations between these events that cannot be explained in terms of a local common cause. That fact is enough to raise the question about non-locality and the compatibility with relativity, regardless of whether one wishes to invoke some ‘holistic’ property of the pair that somehow gives rise to these correlations. In other words, we might simply have to conclude that holism itself is incompatible with relativity. (See Butterfield [1992].)

Second, as discussed earlier in the context of the measurement problem (§5.3.2), entanglement is ubiquitous — even the objects of our everyday experience are likely in entangled states. Can we make sense of the idea that these apparently distinct objects are in fact not independent objects at all, but somehow ‘parts’ of some holistic object? Perhaps, but at the very least, some serious metaphysical work is going to have to be done here, if we are to recover the obvious facts of ordinary experience.

6.2.5 *Relativity Theory*

While some might find non-locality to be contrary to intuition, the real problem that it poses, arguably, is an apparent incompatibility with the theory of relativity. Indeed, supposing that the non-locality of quantum theory were completely and unambiguously compatible with the theory of relativity, it is hard to see how one could have any serious objections to it.

And indeed, on the one hand, one might note that the quantum no-signaling theorem (§6.2.2.2) suggests some kind of compatibility of quantum theory with relativity theory. Hence, there may be, as some have called it, a ‘peaceful co-existence’ between the two theories.⁹⁹ While the quantum correlations seem to imply a failure of locality, and while collapse of the quantum state occurs instantaneously, neither of these facts entails the capacity to send a signal faster than light. More generally, neither of these features of quantum theory entails an *experimental*, or *observational* conflict with relativity. It is sometimes said that quantum theory respects the letter of the law of relativity (but violates the spirit of the law).

Moreover, even the claim that superluminal signaling is incompatible with relativity is controversial. The argument for the incompatibility relies on the idea that such signaling generates inconsistent causal loops because it allows for causal propagation into the past. And some will point out that even superluminal causal processes (for example, the superluminal transmission of matter or energy) is not inconsistent with relativity.

So what *does* relativity demand? One thing seems clear, and is, at least, agreed by the vast majority of physicists: a theory must be Lorentz-invariant to be relativistic. Here, we encounter the real problem. The collapse postulate of standard

⁹⁹The term was introduced by Shimony [1978].

quantum theory is not Lorentz-invariant, nor is it clear how to make it so, while maintaining empirical adequacy. (See Aharonov and Albert [1981].)

There are, of course, relativistic quantum theories. (See 't Hooft, Ch. 7 and Halvorson, Ch. 8, this volume.) Their equations of motion are Lorentz-invariant in the requisite way. However, the collapse postulate, the need for which does not disappear in relativistic quantum theories, is not Lorentz-invariant — it specifies an *instantaneous* collapse.

Moreover, it is far from clear whether the interpretations of quantum theory that most straightforwardly and obviously solve the measurement problem (recall §5.5) can be made Lorentz-invariant. We do not *know* that all hidden-variables theories *must* violate Lorentz-invariance, but the evidence for this claim is very good.

On the other hand, the no-signaling theorem seems to imply that quantum theory itself is observationally consistent with relativity. Hence *any* theory, and in particular any hidden-variables theory, that is observationally indistinguishable from quantum theory, will be observationally consistent with relativity. Hence, while many hidden-variables theories explicitly violate Lorentz-invariance at the level of the evolution of the values of the hidden variables, this violation is not empirically accessible. Another way to put the point is this: these theories require a preferred reference frame (as quantum theory with the collapse postulate itself does), and yet there is no way to determine, experimentally, which is the preferred frame. The question of whether they are thus satisfactory therefore has a lot to do with the status one attaches to Lorentz-invariance.

7 MATHEMATICAL APPENDIX

These notes are intended as quick reminders of several definitions and standard mathematical facts that are assumed in the text. All proofs are omitted.

7.1 Hilbert Spaces

7.1.1 Vector Spaces

A *vector space*, V , is a set that is closed under addition and ‘multiplication by a scalar’ from a field, K . These operations must satisfy, for any $u, v, w \in V$ and $k, k' \in K$: (commutativity) $v + w = w + v$; (vector associativity) $u + (v + w) = (u + v) + w$; (additive identity) $\exists \vec{0} \in V \forall v \in V, v + \vec{0} = v$; (additive inverses) $\forall v \in V, \exists -v \in V, v + (-v) = \vec{0}$, and one generally writes $v - w$ for $v + (-w)$; (scalar associativity) $k(k'v) = (kk')v$; (scalar identity) for $1 \in K$ (the identity in K) $1v = v$; (vector distributivity) $k(v + w) = kv + kw$; (scalar distributivity) $(k + k')v = kv + k'v$. (Notes: commutativity actually follows from the others; additive inverses are (provably) unique; $\forall k \in K, v \in V, k\vec{0} = 0v = \vec{0}$.) In all of the cases that we consider, $K = \mathbb{R}$ or $K = \mathbb{C}$. The set V is used interchangeably to

refer to the entire vector space, and to the underlying set of vectors. Sometimes for clarity one speaks of ‘a vector space over K ’.

7.1.2 Bases and Dimension

Given a set of vectors $\{v_n\} \subseteq V$, any vector of the form $v = \sum_{i=1}^N k_n v_n$ (with $k_n \in K$) is called a *linear combination* of the v_n . The set $\{v_n\}$ is called *linearly independent* if none of the v_n can be written as a linear combination of the others. (In this case, the set may have infinite cardinality, but notice that linear combinations are always finite sums.) Any maximal linearly independent set in V is called a *basis* for V . All such sets provably have the same cardinality, which is called the *dimension* of the space, denoted $\dim V$. All vector spaces, V, V' , with the same dimension are isomorphic. That is, there exists a 1-1 map $m : V \rightarrow V'$ such that, for any $k \in K$ and any $v, w \in V$, $m[k(v + v')] = km(v) + km(v')$. (This last condition makes m a *linear map*, and being 1-1 in addition makes it an isomorphism.)

7.1.3 Inner Product Spaces

An *inner product* on a vector space, V over K , is a map from $V \times V$ to K , denoted $\langle \cdot, \cdot \rangle$, and satisfying, for all $u, v, w \in V$ and all $k \in K$: (non-negativity) $\langle v, w \rangle \geq 0$; (non-degeneracy) $\langle v, v \rangle = 0$ iff $v = \vec{0}$; (sesquilinearity) $\langle u, k(v + w) \rangle = k\langle u, v \rangle + k\langle u, w \rangle$ and $\langle v, w \rangle = \langle w, v \rangle^*$, where $*$ indicates conjugation in K (so, complex conjugation if $K = \mathbb{C}$ and the identity map if $K = \mathbb{R}$). (Note: it follows from the last two properties that $\langle k(v + w), u \rangle = k^*\langle v, u \rangle + k^*\langle w, u \rangle$; hence the name ‘*sesquilinearity*’.) A vector space with an inner product is called an ‘inner product space.’ In an inner product space, w is orthogonal to w' , written $w \perp w'$, if and only if $\langle w, w' \rangle = 0$. In an inner product space, a basis is called ‘orthogonal’ if its elements are mutually orthogonal.

An important fact about inner products is the *Schwarz Inequality*: for any $v, w \in V$, $\langle v, v \rangle \langle w, w \rangle \geq |\langle v, w \rangle|^2$.

7.1.4 Norms and Orthonormal Bases

A *norm* on a vector space, V , over K is a function from V to \mathbb{R} , denoted $\|\cdot\|$, satisfying, for all $v, w \in V$ and $k \in K$: $\|v\| = 0$ iff $v = 0$; $\|kv\| = |k|\|v\|$; and $\|v + w\| \leq \|v\| + \|w\|$. (Note: it follows that $\|v\| \geq 0$ for all $v \in V$.) The norm defines a topology (see §7.5) on V in the obvious way: the open balls are sets of the form $\{x \mid \|x - v\| < r\}$ for some $v \in V$ and $r \in \mathbb{R}$. Another way to put the point is this: a sequence $\{v_n\} \subseteq V$ converges to a vector $v \in V$ ‘in the norm topology’ just in case $\|v_n - v\|$ converges to 0 (in \mathbb{R}). Inner products define a norm by: $\|v\| = \sqrt{\langle v, v \rangle}$. An important fact about normed complex vector spaces in which the norm is derived from an inner product is the ‘polarization identity’:

$$(130) \quad \langle w, v \rangle = \frac{1}{4}(\|w + v\|^2 - \|w - v\|^2 + i\|w + iv\|^2 - i\|w - iv\|^2).$$

In other words, in such spaces, the inner product is also derived from the norm.

In an inner product space, an orthogonal basis (§7.1.3) is called *orthonormal* if its elements all have norm-1. Given any basis for an inner product space, it is possible to construct an orthonormal basis by ‘Gram-Schmidt orthogonalization’, the details of which we will skip. Note that in an inner product space, there is a convenient expression for a given vector, v , in terms of some orthonormal basis, $\{e_n\}$, as: $v = \sum_n \langle e_n, v \rangle e_n$.

7.1.5 Subspaces

A subset, W , of a vector space, V , is a *subspace* if it is a vector space in its own right, under the operations inherited from V . Similar definitions hold for inner product spaces and Hilbert spaces (defined below, §7.1.7). In an inner product space, one subspace, W , is *orthogonal* to another, W' , when, for any $w \in W$ and any $w' \in W'$, $w \perp w'$.

7.1.6 Direct Sums

A *direct sum* of vector spaces V and V' (over the same field, e.g., \mathbb{R} or \mathbb{C}) is the vector space whose elements are taken from the Cartesian product $V \times V'$, with the vector space operations defined component-wise. If $V = V' \oplus V''$, for example, then V' and V'' are disjoint (or, in an inner-product space, orthogonal) subspaces of V , and V is their span.

7.1.7 Banach Space and Hilbert Space

A *Banach space* is a normed vector space that is complete with respect to the norm topology (meaning that the limit of any sequence of vectors is itself contained in the space). A *Hilbert space* is an inner product space that is complete with respect to the norm topology (the norm here being the one given by the inner product). The three ‘classical’ types of Hilbert space are those over the real numbers, complex numbers, and quaternions.

7.1.8 Dual Space

Given a vector space, V , the dual space (sometimes denoted V^* , but the $*$ here is *not* complex conjugation) is the space of linear functionals on V , that is, the space of linear maps from V to \mathbb{R} . When V has topological structure (such as when V is a Hilbert space), we restrict to the continuous linear functionals. The norm of a (continuous) linear functional, ϕ , on V is

$$(131) \quad \|\phi\| = \sup_{v \in V} \{|\phi(v)| \mid \|v\| \leq 1\}.$$

Every finite-dimensional vector space has the same dimension as its dual. The Riesz Representation Theorem states that for any Hilbert space, \mathcal{H} (finite-dimensional or infinite-dimensional), and any continuous linear functional, ϕ , in its dual \mathcal{H}^* ,

there is a unique $v \in \mathcal{H}$ such that, for all $w \in \mathcal{H}$, $\phi(w) = \langle v, w \rangle$. Conversely, each $v \in \mathcal{H}$ obviously generates a continuous linear functional, $\langle v, \cdot \rangle$, on \mathcal{H} . In other words, there is a 1-1 map, $\Phi : \mathcal{H} \rightarrow \mathcal{H}^*$, from a Hilbert space to its dual. Moreover, Φ is an isometry ($\|v\| = \|\Phi(v)\|$) and an ‘anti-isomorphism’ (in particular, $\Phi(kv) = k^*\Phi(v)$ for any $v \in V$ and $k \in \mathbb{C}$). This last property follows from the fact that the inner product is sesquilinear.

7.1.9 Tensor Products

The tensor-product of two Hilbert spaces, \mathcal{H}_1 and \mathcal{H}_2 , both over K , is a third Hilbert space, $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$ over K , constructed from \mathcal{H}_1 and \mathcal{H}_2 as follows. First choose bases $\{e_n\}$ for \mathcal{H}_1 and $\{f_n\}$ for \mathcal{H}_2 . Then form the Cartesian product of $\{e_n\}$ and $\{f_n\}$. This set contains all pairs of the form (e_n, f_m) , and it is stipulated to be a basis for the tensor product space \mathcal{H} . Hence, at this stage, \mathcal{H} consists of all formal linear combinations (over K) of the (e_n, f_m) . Now, let the inner products on \mathcal{H}_1 and \mathcal{H}_2 be denoted $\langle \cdot, \cdot \rangle_1$ and $\langle \cdot, \cdot \rangle_2$. Define the inner product on \mathcal{H} by $\langle v \otimes w, x \otimes y \rangle = \langle v, x \rangle_1 \langle w, y \rangle_2$ for all $v, w \in \mathcal{H}_1$ and $x, y \in \mathcal{H}_2$, and extend to all of $\mathcal{H}_1 \otimes \mathcal{H}_2$ by linearity. Finally, complete $\mathcal{H}_1 \otimes \mathcal{H}_2$ in the norm topology induced by this inner product.

Note that $\dim \mathcal{H} = \dim \mathcal{H}_1 \times \dim \mathcal{H}_2$. One can (if $\dim \mathcal{H}$ is not prime) ‘factorize’ a given Hilbert space into a tensor product, typically in many ways.

7.1.10 Convex Sets and Cones

A subset, X , of a real vector space, V , is *convex* if for any $x, y \in X$, $rx + (1-r)y \in X$ for all r in the real interval $[0, 1]$. In other words, the ‘line segment’ connecting x and y is also in X . A point, x , in a convex set X is an *extreme point* if it is not a convex combination of other points from X . A convex set, X , is a *simplex* if every non-extreme point has a *unique* decomposition as a convex combination of points from X .

A *positive cone* in V is a set $C \subseteq V$ such that $rx \in C$ for all real $r \geq 0$ and all $x \in C$. (A negative cone requires instead $r \leq 0$.) A *convex cone* is a cone that is also convex. Given a set, $S \subseteq V$, one can form the *convex set generated by S* , in the obvious way, by closing under the required condition. Equivalently, one can take the intersection of all convex sets containing S . A similar point holds for cones, and of course for convex cones.

Given a positive cone, C , in V , and an inner product, $\langle \cdot, \cdot \rangle$ on V , the *dual* of C with respect to this inner product is: $C^* = \{y | \langle x, y \rangle \geq 0, x \in C\}$. If $C = C^*$ we say that C is *self-dual*.

A *face*, $F \subseteq X$, of a convex set, X , is a convex subset that is closed under ‘purification’, meaning that for any $v \in F$, if $v = pv_1 + (1-p)v_2$ then v_1 and v_2 are in F as well.

7.2 Operators

7.2.1 Basic Definitions Regarding Operators

A *linear operator*, F , on a vector space, V over K , is a map from V to itself that preserves the linear structure of V , i.e., for any $v, w \in V$ and any $k \in K$: $F(v + w) = F(v) + F(w)$; and $F(kv) = kF(v)$. One normally just writes Fv , and here we refer to F as an ‘operator’ (as we will not be discussing non-linear operators). Two operators, F and G , are said to *commute* if their *commutator*, $[F, G] := FG - GF$ is 0 (i.e., the ‘zero operator’). The *identity operator*, denoted \mathbb{I} , is the operator on V such that $\mathbb{I}v = v, \forall v \in V$.

Given an operator, F , any vector, v , such that $Fv = kv$ for some $k \in K$ is called an *eigenvector* of F , and k is its associated *eigenvalue*. Operators need not have any eigenvectors. The zero vector is not normally counted amongst an operator’s eigenvectors.

An operator, F , on the vector space V is *invertible* just in case there is an operator, G , such that $FG = \mathbb{I}$, where \mathbb{I} is the identity on V . The operator G is denoted F^{-1} . It is unique if it exists.

On a tensor-product space $V = V_1 \otimes V_2$, consider two operators, F on V_1 and G on V_2 . The tensor-product operator $F \otimes G$ can be defined as follows. Choose any bases, $\{e_n\}$ and $\{f_m\}$ for V_1 and V_2 . Define $(F \otimes G)(e_n \otimes f_m) = (Fe_n) \otimes (Gf_m)$, and extend to all of V by linearity.

7.2.2 Boundedness and Continuity

If V has a norm, $\|\cdot\|$, then F is *bounded* if and only if there is some $r \in \mathbb{R}$ such that $\|Fv\| \leq r\|v\|$ for all $v \in V$. When V has a topology (as when it is a Hilbert space), one says that an operator F is *continuous* if and only if it is continuous as a function on V as a topological space. An operator is bounded if and only if it is continuous. Moreover, if the dimension of V is finite, then all operators on V are bounded, i.e., continuous. Note that unbounded operators do not have the entire space as their domain, so that one must keep track of their domains of definition.

7.2.3 Adjoints

When \mathcal{H} is finite-dimensional, we define adjoints and related notions as follows. The *adjoint*, F^* , of an operator, F , on \mathcal{H} satisfies $\langle Fw, v \rangle = \langle w, F^*v \rangle \forall v, w, \in V$. An operator, F , is self-adjoint if $F = F^*$. The proof that adjoints exist is non-trivial.

When \mathcal{H} is infinite-dimensional, we must be slightly more careful, because V might not be in the domain of F , if F is unbounded (see §7.2.2). In the infinite-dimensional case, then, we define the adjoint, F^* , of F as follows. The domain of F^* is all $v \in \mathcal{H}$ such that there is a $v' \in \mathcal{H}$ satisfying $\langle Fw, v \rangle = \langle w, v' \rangle$ for all w in the domain of F . For each such v , define $F^*v = v'$. (One must show that F^* is thus an operator.) Finally, in order to make relevant distinctions in the infinite-dimensional case, we say that an operator, F , is *symmetric* if $\langle Fw, v \rangle = \langle w, Fv \rangle$

for all v, w in the domain of F . We say that F is *self-adjoint* if $F = F^*$. The difference is, provably, that a symmetric but non-self-adjoint operator will have a domain that is a proper subset of the domain of its adjoint.

7.2.4 Normal Operators

A *normal* operator, F , is one that commutes with its adjoint: $FF^* = F^*F$. Given the comments above, it is clear that one must be aware of issues relating to the domains of operators if \mathcal{H} is infinite-dimensional. Clearly all self-adjoint operators are normal, but the converse is not true. Consider $F = 2i\mathbb{I}$.

7.2.5 Projection Operators

An operator, P , is *idempotent* just in case $PP = P$, i.e., $P(Pv) = Pv$ for any $v \in V$. An operator, P , on V is a *projection operator* just in case it is self-adjoint and idempotent. Each projection operator, P , corresponds to a closed subspace, namely, the subspace of vectors, v , for which $Pv = v$. Note that $PQ = 0$ if and only if the corresponding subspaces are orthogonal.

7.2.6 Unitary Operators

An *automorphism* of a vector space, V , is a map from V to itself that ‘preserves the structure of V ’, and in particular the linear, inner-product, and topological structures (the latter two if they exist in V). Let U be a (linear) operator on the Hilbert space \mathcal{H} such that: U is invertible (hence U is 1-1); and U preserves inner products (i.e., for any $v, w \in \mathcal{H}$, $\langle Uw, Uv \rangle = \langle w, v \rangle$). Such an operator is called ‘unitary’, and clearly implements an automorphism of \mathcal{H} . In that case, of course U also preserves norms, i.e., $\|Uv\| = \|v\|$ for all v .

It is readily shown that for any unitary operator, U , $U^* = U^{-1}$. Conversely, any invertible linear operator with that property is unitary. (If, instead, $\langle Uw, Uv \rangle = \langle w, v \rangle^*$ and $U(kv) = k^*Uv$, then U is *anti-unitary*.)

7.3 The Hilbert Space \mathbb{C}^2

The space of complex column-vectors with 2 components is denoted \mathbb{C}^2 . The elements of \mathbb{C}^2 are added component-wise:

$$(132) \quad \begin{pmatrix} a \\ b \end{pmatrix} + \begin{pmatrix} c \\ d \end{pmatrix} = \begin{pmatrix} a + c \\ b + d \end{pmatrix}.$$

The (linear) operators on this space can be represented by 2×2 complex matrices. A matrix operates on a vector according to the rule:

$$(133) \quad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} ax + by \\ cx + dy \end{pmatrix}.$$

The product of two matrices (which can in fact be derived from the rule above) is:

$$(134) \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} e & f \\ g & h \end{pmatrix} = \begin{pmatrix} ae + bg & ce + dg \\ af + bh & cf + dh \end{pmatrix}.$$

The inner product on the space is given by

$$(135) \left\langle \begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} c \\ d \end{pmatrix} \right\rangle = a^*c + b^*d.$$

The norm, or ‘length’, of the vector with components a, b is then just $\sqrt{a^*a + b^*b}$. (Compare this expression with the Euclidean length of a vector in the real vector space \mathbb{R}^2 .)

7.3.1 The Pauli Matrices

The Pauli Matrices (operators on \mathbb{C}^2) are:

$$(136) \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

These matrices have many nice properties, which the reader might wish to verify. For example:

$$(137) \begin{aligned} \sigma_x^2 &= \sigma_y^2 = \sigma_z^2 = \mathbb{I} \\ \text{Tr}[\sigma_x] &= \text{Tr}[\sigma_y] = \text{Tr}[\sigma_z] = 0 \\ \sigma_x\sigma_y &= i\sigma_z \\ [\sigma_x, \sigma_y] &= 2i\sigma_z. \end{aligned}$$

The last two properties generalize: they remain true under cyclic permutations of the indices.

One often sees the expression $\vec{\sigma}$, which is to be understood as a ‘vector’ whose components are the three Pauli matrices, so that, for example, $\vec{r} \cdot \vec{\sigma}$ is a shorthand for a linear combination of the three Pauli matrices, with the coefficients given by the components of (the real vector) \vec{r} .

7.4 Posets and Lattices

7.4.1 Posets

A *partially ordered set* (normally, *poset*) is a set, L , together with a relation, \leq , that obeys, for all $a, b, c \in L$: (reflexivity) $a \leq a$; (anti-symmetry) if $a \leq b$ and $b \leq a$ then $a = b$; and (transitivity) if $a \leq b$ and $b \leq c$ then $a \leq c$. The relation \leq is called a *partial order* on L . Note, in particular, that in general there will be $a, b \in L$ such that neither $a \leq b$ nor $b \leq a$. (If there are no such a and b in L , then \leq is a total order on L .) As elsewhere, I let L denote both the poset itself as well as its underlying set.

7.4.2 Lattices

Let L be a poset. Define the *join* of two elements, $a, b \in L$ as the least upper bound of a and b , i.e., the smallest (under \leq) c such that $a \leq c$ and $b \leq c$. The join of a and b (which need not exist in general) is denoted $a \vee b$. Define the *meet* of two elements, $a, b \in L$ as the greatest lower bound of a and b , i.e., the largest (under \leq) c such that $c \leq a$ and $c \leq b$. The meet of a and b (which need not exist in general) is denoted $a \wedge b$. A poset in which every pair (hence every finite set) of elements has a meet and a join is called a *lattice*.

A lattice, L , is *complemented* if every $a \in L$ has a *complement*, $a^\perp \in L$, satisfying: $a \wedge a^\perp = 0$; and $a \vee a^\perp = 1$. The operation $^\perp$ is in this case a *complement*. A complemented lattice L is called *orthocomplemented*, or an *ortholattice*, if for all $a, b \in L$: $a \leq b$ implies $b^\perp \leq a^\perp$; and $a^{\perp\perp} = a$. The operation $^\perp$ is in this case an *orthocomplement*.

7.4.3 Distributivity

A lattice, L , is *distributive* if, for all $a, b, c \in L$: $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$, and similarly, swapping \wedge with \vee . In general, lattices are not distributive. The *center*, $Z(L)$, of a lattice L is the set of all $z \in L$ such that, for any $p \in L$, $p = (p \wedge z) \vee (p \wedge z^\perp)$. $Z(L)$ is a distributive sublattice of L .

7.4.4 Direct Products and Reducibility

Let L_1 and L_2 be ortholattices. Then we can form a third ortholattice, L , which is the ‘direct product’ of L_1 and L_2 . As a set, L is the direct (Cartesian) product of L_1 and L_2 as sets. Then define, for $a, b \in L$ where $a = (a_1, a_2)$ and $b = (b_1, b_2)$, $a \leq b$ just in case $a_1 \leq b_1$ and $a_2 \leq b_2$. Meets, joins, and orthocomplements are similarly (therefore) defined componentwise. An ortholattice, L is ‘irreducible’ if it is not isomorphic to the direct product of non-trivial ortholattices.

7.4.5 Atomicity and the Covering Property

The ‘top’ (or ‘identity’) of a lattice (if it exists) is the element $\mathbb{I} \in L$ such that $a \leq \mathbb{I}$ for all $a \in L$. The ‘bottom’ (or ‘zero’) of a lattice (if it exists) is the element $0 \in L$ such that $0 \leq a$ for all $a \in L$. An *atom* in a poset, L , is a non-zero element $a \in L$ such that, for any $b \in L$, if $b \leq a$ then either $b = 0$ or $b = a$. A poset is *atomic* if every non-zero element contains (under \leq) an atom. Finally, a lattice is *complete* if every collection of elements from L has a meet and join. In a complete atomic lattice, every element (apart from 0) is either an atom or the join of atoms.

Finally, we say that a lattice L has the *covering property* if, for every atom $a \in L$ and any $b \in L$ where $a \wedge b = 0$, $a \vee b$ ‘covers’ b ; that is, there is no element strictly between b and $a \vee b$.

7.5 Topology and Measure

7.5.1 Topological Spaces

A *topological space* is a set, S , together with a collection, T , of subsets of S satisfying: $\emptyset, S \in T$ (with \emptyset the null set); the union of any collection of sets in T is in T ; the intersection of any pair of sets in T is in T . T is the *topology* of S . The elements of T are the *open sets*, and their complements (in S) are the *closed sets*. A function, f , from one topological space to another is *continuous* if the inverse image under f of every open set is again an open set.

A *base*, B , for a topology, T , is a collection of open sets in T such that every open set in T can be written as a union of elements of B . A common example involves spaces (such as vector spaces) with a norm, where one can define a base as the collection of ‘open balls’, that is, the collection of sets of the form $\{x \mid \|x - y\| < \epsilon\}$ where x and y are points (e.g., vectors) in the space, and ϵ is a real number.

A topological space is *compact* if every sequence of points has a subsequence that converges to some point in the space. The space is *locally compact* if, roughly, each point in the space has a neighborhood that is compact — roughly, each small part of the space ‘looks like’ a small part of a compact space.

Given two topological spaces, S_1 and S_2 , we can form their Cartesian product, $S_1 \times S_2$ as sets. We then define the *product topology* on this Cartesian product as follows. Let B_1 and B_2 be bases for the topologies on S_1 and S_2 respectively. The product topology on $S_1 \times S_2$ is the topology whose base is the Cartesian product $B_1 \times B_2$. (The definition of the product topology on a Cartesian product of more than two spaces is more convoluted. We will not need it here.)

7.5.2 Manifolds

A *manifold* is a topological space that is ‘locally Euclidean’, meaning that there is a neighborhood around each point that is topologically the same as the open unit ball in \mathbb{R}^n (for some n — and n is then the *dimension* of the manifold). An open set, S , of the manifold M , together with a homeomorphism between S and an open set of \mathbb{R}^n is called a *coordinate chart*. A collection of charts that covers M is an *atlas*. Now consider the region, S , in which two charts overlap. We thus have two distinct maps from S to \mathbb{R}^n , which defines a map, μ , from a subset of \mathbb{R}^n (the range of the first chart, applied to S) to some other subset of \mathbb{R}^n (the range of the second chart, applied to S). If all of the μ generated by all overlapping charts in the atlas are infinitely differentiable, then the manifold is a *smooth manifold*.

7.5.3 Weak Operator Topology

The weak operator topology over the operators on a Hilbert space, \mathcal{H} is the weakest topology on the set, $\mathcal{B}(\mathcal{H})$, of bounded operators on \mathcal{H} , such that the map $F \mapsto \langle w, Fv \rangle$ is continuous for any vectors $v, w \in \mathcal{H}$ and any $F \in \mathcal{B}(\mathcal{H})$. In the weak operator topology, a sequence, $\{F_n\}$, of operators converges to the operator F just in case $|\langle w, F_n v \rangle - \langle w, F v \rangle|$ converges to 0 for every $v, w \in \mathcal{H}$.

7.5.4 Lebesgue Measure

On the real line, \mathbb{R} , the *Lebesgue measure* is the natural extension of the usual measure of distance (size of intervals) to more complicated sets of points. For example, given any open set, S , that is the union of disjoint intervals, the Lebesgue measure of S is the sum of the size of the intervals. Any countable union of individual points in \mathbb{R} has Lebesgue-measure zero. The measure is extended to volumes in \mathbb{R}^3 in the obvious way.

7.5.5 Borel Sets

Borel sets of real numbers are definable as follows. Given some set, S , a σ -algebra over S is a family of subsets of S closed under complement, countable union and countable intersection. The *Borel algebra* over \mathbb{R} is the smallest σ -algebra containing the open sets of \mathbb{R} . (One must show that there is indeed a smallest.) A *Borel set* of real numbers is an element of the Borel algebra over \mathbb{R} . Note that not every subset of real numbers is a Borel set, though the ones that are not are somewhat exotic. All open and closed sets are Borel. The importance of Borel algebras (hence Borel sets) lies in the fact that certain measure-theoretic results apply only to them. On the other hand, in many cases one can extend the important results and definitions to a wider class of sets, for example, all sets that are the image of a Borel set under a continuous function. However, we shall not continue to make note of such points.

7.5.6 Probability Measures

Let X be a set (a ‘sample space’) of basic events and \mathcal{A} a σ -algebra over X . The (Kolmogorovian) axioms of probability theory may then be stated as follows. Let $p : \mathcal{A} \rightarrow [0, 1]$ be a map from \mathcal{A} to the interval $[0, 1]$. This p is a *probability measure* on \mathcal{A} just in case: (normalization) $p(X) = 1$; (negation) $p(E') = 1 - p(E)$ for any $E \in \mathcal{A}$; and (additivity) for any countable disjoint sequence $\{E_k\}$ of elements of \mathcal{A} , $p(\cup_k E_k) = \sum_k p(E_k)$.

7.6 Groups

7.6.1 Groups and Homomorphisms

A *group* is a non-empty set, \mathcal{G} , with a binary operation, $*$ (called the ‘product’) on that set satisfying: (associativity) for all $a, b, c \in \mathcal{G}$, $(a * b) * c = a * (b * c)$; (identity) there is an element, $e \in \mathcal{G}$, such that for any $a \in \mathcal{G}$, $e * a = a * e = a$; (inverse) for all $a \in \mathcal{G}$, there exists $b \in \mathcal{G}$ such that $a * b = b * a = e$ (this b is usually denoted a^{-1}); (closure) for all $a, b \in \mathcal{G}$, $a * b \in \mathcal{G}$. A structure satisfying all of these properties except for the existence of inverses is called a ‘semi-group’. A map, $m : \mathcal{G} \rightarrow \mathcal{G}'$, from one group, \mathcal{G} , to another, \mathcal{G}' , is a *homomorphism* if, for all $a, b \in \mathcal{G}$, $m(a * b) = m(a) * m(b)$. (Note that the product on the left-hand side is taken in \mathcal{G} , while the product on the right-hand side is taken in \mathcal{G}' .) It follows

that m preserves inverses and maps the identity in \mathcal{G} to the identity in \mathcal{G}' . The map m is an isomorphism if it is a homomorphism and it is 1-1. An isomorphism from \mathcal{G} to itself is an *automorphism* of \mathcal{G} . Normally the group product operator $*$ is left implicit; hence we will from now on write $a * b$ as ab and so on.

7.6.2 Subgroups and Products

Given a group, \mathcal{G} , the subgroup \mathcal{H} is *normal* if and only if $g\mathcal{H}g^{-1} \subset \mathcal{H}$ for any $g \in \mathcal{G}$ (where $g\mathcal{H}g^{-1}$ is the set $\{ghg^{-1} | h \in \mathcal{H}\}$). Given two groups, \mathcal{H} and \mathcal{K} , the group \mathcal{G} is their *direct product* if and only if: (i) \mathcal{H} and \mathcal{K} are (isomorphic to) normal subgroups of \mathcal{G} ; (ii) $\mathcal{H} \cap \mathcal{K} = e$, the identity in \mathcal{G} ; and (iii) as a set, \mathcal{G} is (isomorphic to) $\{hk : h \in \mathcal{H}, k \in \mathcal{K}\}$. We write $\mathcal{G} = \mathcal{H} \times \mathcal{K}$. The usual way to construct a direct product of groups is to let \mathcal{G} as a set be the Cartesian product of \mathcal{H} and \mathcal{K} as sets, and define $(h, k)(h', k') = (hh', kk')$ for all $h, h' \in \mathcal{H}$ and $k, k' \in \mathcal{K}$. (Notice in this case that \mathcal{H} is isomorphic to the subgroup of elements of the form (h, e) for any $h \in \mathcal{H}$ and e the identity in \mathcal{G} , and similarly for \mathcal{K} . Both \mathcal{H} and \mathcal{K} in this case are also normal subgroups.) If only \mathcal{K} is a normal subgroup of \mathcal{G} , while \mathcal{H} is a non-normal subgroup, then \mathcal{G} is the *semi-direct product* of \mathcal{H} by \mathcal{K} , and we write $\mathcal{G} = \mathcal{H} \rtimes \mathcal{K}$.

7.6.3 Cosets and Quotients

Let \mathcal{H} be a subgroup of \mathcal{G} . We define the (*left*) *cosets of \mathcal{H} in \mathcal{G}* as the set $\{g\mathcal{H} | g \in \mathcal{G}\}$, where $g\mathcal{H} = \{gh | h \in \mathcal{H}\}$. (Note that for some of the $g, g' \in \mathcal{G}$ we will have $g\mathcal{H} = g'\mathcal{H}$.) The left cosets of \mathcal{H} in \mathcal{G} partition \mathcal{G} . They also themselves form a group, with the multiplication rule $(g\mathcal{H})(g'\mathcal{H}) = (gg')\mathcal{H}$. This group is called the *quotient of \mathcal{G} by \mathcal{H}* , typically written \mathcal{G}/\mathcal{H} .

7.6.4 Continuously Parametrized Groups

To say that the group, \mathcal{G} , is *continuously parametrized by a* means, in particular, that the group as a set is indexed by a , and: $g_0 = I$ (the identity on \mathcal{G}); $g_{a+b} = g_a g_b$; $\lim_{n \rightarrow \infty} a_n = b$ implies $\lim_{n \rightarrow \infty} g_{a_n} = g_b$. In the third condition, the limit on the right requires that the group be a *topological group*, i.e., the group is also a topological space.

7.6.5 Lie Groups

A *Lie group* is a smooth manifold (§7.5.2) that is also a group, where the group operations of multiplication and inversion are continuous maps on the manifold. (In fact, Lie groups are often defined as analytic manifolds, but we will not bother with that point here.)

7.6.6 Vector Space Representations

It is a fundamental theorem of group theory that every group, \mathcal{G} , is (isomorphic to) a subgroup of the group of permutations on some set. Another common type of representation of \mathcal{G} is a *vector space representation*, a group-homomorphism from \mathcal{G} to $\text{GL}(V)$, the ‘general linear group’ of transformations of the vector space V , that is, the group of invertible linear operators on V .

A representation $m : \mathcal{G} \rightarrow \text{GL}(V)$ of the group \mathcal{G} is *faithful* just in case m is 1-1. Non-faithful representations ignore structure in the represented group. Representations can also *introduce* structure, in the following sense. A proper subspace W of V that is invariant under the group (that is, the group is an automorphism of W) carries a ‘subrepresentation’ of \mathcal{G} in the sense that the restriction of the representation of \mathcal{G} on V to W is itself a representation of \mathcal{G} . When a (proper) subrepresentation of a representation, m , of \mathcal{G} exists, m is called ‘reducible’. Otherwise, it is ‘irreducible’. If a representation is reducible, then the vector space on which it is represented is in a sense ‘bigger than it needs to be’ to represent the group.

7.6.7 Group Action

For any set, S , a *group action* of \mathcal{G} on S is a map $\mu : \mathcal{G} \times S \rightarrow S$ satisfying: (a) $\mu(g, \mu(h, s)) = \mu(gh, s)$ for all $g, h \in \mathcal{G}$ and $s \in S$; and (b) $\mu(e, s) = s$ for all $s \in S$ (where e is the identity in \mathcal{G}). $\mu(g, s)$ is often written $\mu_g s$. Each μ_g is in fact a bijection on S , so that we may also define a group action as a group homomorphism from \mathcal{G} to the group of bijections on S . Sometimes μ_g is called the ‘action of g on S ’. A group action, μ , of \mathcal{G} on S is *transitive* if and only if for any $s, t \in S$, there exists $g \in \mathcal{G}$ such that $\mu_g s = t$. If both \mathcal{G} and S have a topological structure, then the action of \mathcal{G} on S is *continuous* if the map μ is continuous with respect to the product topology (§7.5.1) of $\mathcal{G} \times S$.

7.6.8 Unitary Representations

Given a Hilbert space, \mathcal{H} , any unitary operator, U , on \mathcal{H} implements an automorphism of $\text{GL}(\mathcal{H})$. In particular, map $\text{GL}(\mathcal{H})$ to itself via $F \mapsto U^{-1}FU$ for every $F \in \text{GL}(\mathcal{H})$. Note that, in particular, $U^{-1}FGU = (U^{-1}FU)(U^{-1}GU)$ and $U^{-1}F^{-1}U = (U^{-1}FU)^{-1}$ (assuming F is invertible), for any $F, G \in \text{GL}(\mathcal{H})$ (using the fact that $(AB)^{-1} = B^{-1}A^{-1}$ for any invertible operators A and B). Moreover, note that this map is 1-1: for any operator F , there is a unique operator G such that $U^{-1}GU = F$, namely, $G = UFU^{-1}$.

Note also that every unitary operator is an element of $\text{GL}(\mathcal{H})$. Hence we have here an example of a general construction, the *group inner isomorphism*, by which one maps a group, \mathcal{G} , to itself via $h \mapsto ghg^{-1}$ for some $g \in \mathcal{G}$ and all $h \in \mathcal{G}$.

7.6.9 Induced Representations

Given a group, \mathcal{G} , and a subgroup, \mathcal{H} , a representation of \mathcal{H} on a vector space, W , ‘induces’ a representation of \mathcal{G} , in the following sense. One can construct, from the representation of \mathcal{H} on W , a vector space, V , which is in fact the direct sum of copies of W , and a representation of \mathcal{G} on V . Each of the copies of W inside V carries a representation of \mathcal{H} .

Here is a rough description of how the construction goes. Let σ be a representation of \mathcal{H} on W . The general idea behind constructing the representation of \mathcal{G} induced by σ is to construct a vector space V that is the direct sum of copies of W , i.e., $V = \bigoplus_n W_n$, where each W_n is a copy of W , and each copy of W corresponds to an element of \mathcal{G}/\mathcal{H} . The representation, ρ , of \mathcal{G} induced by σ is defined as follows. Fix a representative, g_n , for each coset, n , of \mathcal{G}/\mathcal{H} . Note that for any $g \in \mathcal{G}$, $gg_n = g_m h$ for some $h \in \mathcal{H}$ and some $m \in \mathcal{G}/\mathcal{H}$. Let w_n be an arbitrary vector from W_n , corresponding (under some isomorphism) to $w \in W$. (Note that if we define an operator on a basis for each of the W_n then we have defined it for all of V , by linearity.) Define the $\rho(g)$ (for any $g \in \mathcal{G}$) by $\rho(g)w_n = (\sigma(h)w)_m$, where h and m are given as above. The expression on the right should be read as ‘let $\sigma(h)$ act on $w_n \in W_n$ in the same way that it acts on $w \in W$, then map the result to the corresponding vector in W_m .’ Note that this entire prescription presupposes a set of isomorphisms between W and the W_n . Finally, it can be shown that none of the above depends on the choice of representatives g_n , in the sense that a different choice produces an isomorphic representation.

ACKNOWLEDGEMENTS

Thanks to Jeremy Butterfield and John Earman for their comments and suggestions. Thanks also to the participants of a workshop at the University of Pittsburgh in November, 2004, for helpful comments.

BIBLIOGRAPHY

- [Aharonov and Albert, 1981] Yakir Aharonov and David Albert. Can we make sense of the measurement process in relativistic quantum mechanics? *Physical Review Letters*, 24:359–370, 1981.
- [Aharonov and Botero, 2005] Yakir Aharonov and Alonso Botero. Quantum averages of weak values. *Physical Review A*, 72:052111, 2005.
- [Aharonov and Kaufherr, 1984] Yakir Aharonov and M. Kaufherr. Quantum frames of reference. *Physical Review D*, 30:368–385, 1984.
- [Aharonov *et al.*, 1987] Yakir Aharonov, David Albert, A. Casher, and Lev Vaidman. Surprising quantum effects. *Physics Letters A*, 124:199–203, 1987.
- [Aharonov *et al.*, 1993] Yakir Aharonov, Jeeva Anandan, and Lev Vaidman. Meaning of the wavefunction. *Physical Review A*, 47:4616–4626, 1993.
- [Albert and Loewer, 1988] D. Albert and B. Loewer. Interpreting the many worlds interpretation. *Synthese*, 77:195–213, 1988.
- [Alfsen and Shultz, 2003] Erik M. Alfsen and Frederic W. Shultz. *Geometry of State Spaces of Operator Algebras*. Birkhäuser, Boston, 2003.

- [Atmanspacher, 2004] Harald Atmanspacher. Quantum approaches to consciousness. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2004. <http://plato.stanford.edu/archives/win2004/entries/qt-consciousness/>.
- [Bacciagaluppi and Dickson, 1999] Guido Bacciagaluppi and Michael Dickson. Dynamics for modal interpretations. *Foundations of Physics*, 29:1165–1201, 1999.
- [Bacciagaluppi and Hemmo, 1996] Guido Bacciagaluppi and Meir Hemmo. Modal interpretations, decoherence, and measurements. *Studies in History and Philosophy of Modern Physics*, 27:239–277, 1996.
- [Bacciagaluppi, 2000] Guido Bacciagaluppi. Delocalized properties in the modal interpretation of a continuous model of decoherence. *Foundations of Physics*, 30:1431–1444, 2000.
- [Ballentine, 1970] L. Ballentine. The statistical interpretation of quantum mechanics. *Reviews of Modern Physics*, 42:358–381, 1970.
- [Barbour, 1989] J. Barbour. *The Discovery of Dynamics: A Study from a Machian Point of View of the Discovery and the Structure of Dynamical Theories*. Oxford University Press, Oxford, 1989.
- [Barrett, 1999] J. Barrett. *The Quantum Mechanics of Minds and Worlds*. Oxford University Press, Oxford, 1999.
- [Bedard, 1999] K. Bedard. Material objects in Bohm’s interpretation. *Philosophy of Science*, 66:221–242, 1999.
- [Bell, 1964] John Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.
- [Bell, 1990] John Bell. Against measurement. In A. Miller, editor, *Sixty-Two Years of Uncertainty*, pages 17–31. Plenum, New York, 1990.
- [Beller, 1999] Mara Beller. *Quantum Dialogue*. University of Chicago Press, Chicago, 1999.
- [Birkhoff and von Neumann, 1936] G. Birkhoff and J. von Neumann. The logic of quantum mechanics. *Annals of Mathematics*, 37:823–843, 1936.
- [Bohm, 1951] David Bohm. *Quantum Theory*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1951.
- [Bohm, 1952] David Bohm. A suggested interpretation of the quantum theory in terms of ‘hidden’ variables, I and II. *Physical Review*, 85:166–193, 1952.
- [Bohr, 1935] Niels Bohr. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 48:696–702, 1935.
- [Bracken, 2003] A. J. Bracken. Quantum mechanics as an approximation to classical mechanics in Hilbert space. *Journal of Physics A*, 36:L329–L335, 2003.
- [Brown and Pooley, 2002] Harvey Brown and Oliver Pooley. Relationalism rehabilitated? I: Classical mechanics. *British Journal for the Philosophy of Science*, 53:183–204, 2002.
- [Brown, 1986] Harvey Brown. The insolubility proof of the quantum measurement problem. *Foundations of Physics*, 16:857–870, 1986.
- [Bub and Clifton, 1996] Jeffrey Bub and Rob Clifton. A uniqueness theorem for interpretations of quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 27B:181–219, 1996.
- [Bub, 1974] Jeffrey Bub. *The Interpretation of Quantum Mechanics*. D. Reidel, Dordrecht, 1974.
- [Bub, 1977] Jeffrey Bub. Von Neumann’s projection postulate as a possibility conditionalization rule in quantum mechanics. *Journal of Philosophical Logic*, 6:381–390, 1977.
- [Bub, 1997] Jeffrey Bub. *Interpreting the Quantum World*. Cambridge University Press, Cambridge, 1997.
- [Busch *et al.*, 1995] P. Busch, M. Grabowski, and P. Lahti. *Operational Quantum Physics*. Springer-Verlag, Berlin, 1995.
- [Busch, 2003] Paul Busch. Quantum states and generalized observables: A simple proof of Gleason’s theorem. *Physical Review Letters*, 91:1–4, 2003.
- [Butterfield, 1992] Jeremy Butterfield. Bell’s theorem: What it takes. *British Journal for the Philosophy of Science*, 43:41–83, 1992.
- [Butterfield, 1993] Jeremy Butterfield. Interpretation and identity in quantum theory. *Studies in History and Philosophy of Science*, 14:443–478, 1993.
- [Butterfield, 1995] Jeremy Butterfield. Worlds, minds, and quanta. *Aristotelian Society Supplementary Volume*, 69:113–158, 1995.
- [Cattaneo, 1979] U. Cattaneo. Mackey’s imprimitivity theorem. *Commentari Mathematici Helvetici*, 54:629–641, 1979.
- [Clauser and Shimony, 1978] J. Clauser and A. Shimony. Bell’s theorem: Experimental tests and implications. *Reports on Progress in Physics*, 41:1881–1927, 1978.

- [Clifton, 1995] Rob Clifton. Independently motivating the Kochen-Dieks modal interpretation of quantum mechanics. *Commentari Mathematici Helvetici*, 46:33–57, 1995.
- [Clifton, 2000] Rob Clifton. The modal interpretation of algebraic quantum field theory. *Physics Letters A*, 271:161–177, 2000.
- [Cohen and Svetlichny, 1987] D. Cohen and G. Svetlichny. Minimal supports in quantum logics. *International Journal of Theoretical Physics*, 27:435–450, 1987.
- [Cohen-Tannoudji *et al.*, 1977] C. Cohen-Tannoudji, D. Bernard, and F. Laloë. *Quantum Mechanics*. John Wiley & Sons, New York, 1977.
- [Collier, 1999] J. D. Collier. Causation is the transfer of information. In H. Sankey, editor, *Causation, Natural Laws and Explanation*, pages 279–331. Kluwer, Dordrecht, 1999.
- [Cushing and McMullin, 1989] J. T. Cushing and E. McMullin, editors. *Philosophical Consequences of Quantum Theory*. University of Notre Dame Press, Notre Dame, IN, 1989.
- [Cushing *et al.*, 1996] J. T. Cushing, A. Fine, and S. Goldstein, editors. *Bohmian Mechanics and Quantum Theory: An Appraisal*, volume 184 of *Boston Studies in the Philosophy of Science*. Kluwer Academic Publishers, Boston, 1996.
- [Cushing, 1994] J. T. Cushing. *Quantum Mechanics: Historical Contingency and the Copenhagen Hegemony*. University of Chicago Press, Chicago, 1994.
- [dalla Chiara *et al.*, 2004] M. dalla Chiara, R. Giuntini, and R. Ghrechie. *Reasoning in Quantum Theory: Sharp and Unsharp Quantum Logics*. Kluwer Academic Publishers, Dordrecht, 2004.
- [de Broglie, 1924] Louis de Broglie. *Recherches sur la Théorie des Quanta*. PhD thesis, Faculté des Sciences de Paris, 1924. Appears in *Annales de Physique*, series 10, vol III, pp. 22–128, 1925.
- [de Broglie, 1927] Louis de Broglie. La nouvelle dynamique des quanta. In H. Lorentz, editor, *Rapports et Discussions du Cinquième Conseil de Physique Solway*, pages 105–141, Paris, 1927. Gauthier-Villars.
- [d’Espagnat, 1971] Bernard d’Espagnat. *The Conceptual Foundations of Quantum Mechanics*. W.A. Benjamin, Inc., Reading, MA, 1971.
- [Dewdney *et al.*, 1986] C. Dewdney, P. Holland, and A. Kyprianidis. What happens in a spin measurement? *Physics Letters A*, 119:259–267, 1986.
- [DeWitt and Graham, 1973] B. S. DeWitt and N. Graham, editors. *The Many-Worlds Interpretation of Quantum Mechanics*. Princeton University Press, Princeton, 1973.
- [Dickson, 1998] Michael Dickson. *Quantum Chance and Nonlocality*. Cambridge University Press, Cambridge, 1998.
- [Dickson, 2000] Michael Dickson. Are there material objects in Bohm’s theory? *Philosophy of Science*, 67:704–710, 2000.
- [Dickson, 2001] Michael Dickson. Quantum logic is alive \wedge (it is true \vee it is false). *Philosophy of Science*, 68:S274–S287, 2001.
- [Dickson, 2002a] Michael Dickson. Bohr on Bell: A proposed reading of Bohr and its implications for Bell’s theorem. In J. Butterfield and T. Placek, editors, *Proceedings of the NATO Advanced Research Workshop on Modality, Probability and Bell’s Theorem*, Amsterdam, 2002. IOS Press.
- [Dickson, 2002b] Michael Dickson. The EPR experiment: A prelude to Bohr’s reply to EPR. In Michael Heidelberger and Friedrich Stadler, editors, *History of Philosophy of Science—New Trends and Perspectives*, volume 9 of *Institute Vienna Circle Yearbook*, pages 263–276, Dordrecht, 2002. Kluwer.
- [Dickson, 2004a] Michael Dickson. Quantum reference frames in the context of EPR. *Philosophy of Science*, 71:655–668, 2004.
- [Dickson, 2004b] Michael Dickson. A view from nowhere: Quantum reference frames and uncertainty. *Studies in History and Philosophy of Modern Physics*, 35:195–220, 2004.
- [Dieks and Vermaas, 1998] Denis Dieks and Pieter Vermaas, editors. *The Modal Interpretation of Quantum Mechanics*. The Western Ontario Series in Philosophy of Science. Kluwer Academic Publishers, Dordrecht, 1998.
- [Dieks, 1988] Denis Dieks. The formalism of quantum theory: An objective description fo reality? *Annalen der Physik*, 7:174–190, 1988.
- [Diosi, 1992] L. Diosi. Quantum measurement and gravity for each other. In P. Cvitanovic, I. Pervical, and A. Wirzba, editors, *Quantum Chaos—Quantum Measurement*. Kluwer, Dordrecht, 1992.

- [Dirac, 1930] P. A. M. Dirac, editor. *The Principles of Quantum Mechanics*. Clarendon Press, Oxford, 1st edition, 1930.
- [DiSalle, 1990] Robert DiSalle. Conventionalism and the origins of the inertial frame concept. In Arthur Fine, Micky Forbes, and Linda Wessels, editors, *PSA 1990*, East Lansing, MI, 1990. The Philosophy of Science Association.
- [DiSalle, 2002] Robert DiSalle. Space and time: Inertial frames. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2002. <http://plato.stanford.edu/archives/sum2002/entries/spacetime-iframes>.
- [Dixmier, 1958] J. Dixmier. Sur la relation $i(pq-qp)=1$. *Compositio Mathematica*, 13:263–270, 1958.
- [Donald, 1990] M. Donald. Quantum theory and the brain. *Proceedings of the Royal Society of London*, A427:43–93, 1990.
- [Dummett, 1976] Michael Dummett. Is logic empirical? In H. Lewis, editor, *Contemporary British Philosophy*. Allen and Unwin, London, 1976.
- [Earman and Ruetsche, 2006] John Earman and Laura Ruetsche. Relativistic invariance and modal interpretations. *Philosophy of Science*, forthcoming, 2006.
- [Einstein *et al.*, 1935] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, 1935.
- [Einstein, 1905] Albert Einstein. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 17:132–149, 1905.
- [Everett, 1957] H. Everett. Relative state formulations of quantum mechanics. *Reviews of Modern Physics*, 29:454–462, 1957.
- [Fine, 1973] Arthur Fine. Probability and the interpretation of quantum mechanics. *British Journal for the Philosophy of Science*, 24:1–37, 1973.
- [Fine, 1986] Arthur Fine. *The Shaky Game: Einstein, Realism and the Quantum Theory*. University of Chicago Press, Chicago, 1986.
- [Fine, 1991] Arthur Fine. Inequalities for nonideal correlation experiments. *Foundations of Physics*, 21:365–378, 1991.
- [Finkelstein, 1962] D. Finkelstein. The logic of quantum physics. *Transactions of the New York Academy of Sciences*, 25:621–637, 1962.
- [Finkelstein, 1969] D. Finkelstein. Matter, space, and logic. In M. Wartofsky and R. Cohen, editors, *Boston Studies in the Philosophy of Science*, volume 5, pages 199–215. Humanities Press, New York, 1969.
- [Friedman and Putnam, 1978] Michael Friedman and Hilary Putnam. Quantum logic, conditional probability, and interference. *Dialectica*, 32:305–315, 1978.
- [Ghriardi *et al.*, 1986] G.-C. Ghriardi, A. Rimini, and T. Weber. Unified dynamics for microscopic and macroscopic systems. *Physical Review D*, 34:470–491, 1986.
- [Ghriardi *et al.*, 1990] G.-C. Ghriardi, R. Grassi, and P. Pearle. Relativistic dynamic reduction models—General framework and examples. *Foundations of Physics*, 20:1271–1316, 1990.
- [Gibbons, 1987] P. Gibbons. *Particles and Paradoxes: The Limits of Quantum Logic*. Cambridge University Press, Cambridge, 1987.
- [Gisin, 1984] N. Gisin. Quantum measurement and stochastic processes. *Physical Review Letters*, 52:1657–1660, 1984.
- [Gleason, 1957] A. Gleason. Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6:885–893, 1957.
- [Greenberger *et al.*, 1989] D. M. Greenberger, M. A. Horne, and A. Zeilinger. Going beyond Bell's theorem. In M. Kafatos, editor, *Bell's Theorem, Quantum Theory, and Conceptions of the Universe*, pages 69–72. Kluwer, Dordrecht, 1989.
- [Groenewold, 1946] H. J. Groenewold. On the principles of elementary quantum mechanics. *Physica*, 12:405–460, 1946.
- [Haag, 1992] Rudolf Haag. *Local Quantum Physics*. Springer-Verlag, Berlin, 1992.
- [Halvorson and Clifton, 1999] Hans Halvorson and Rob Clifton. Maximal beable subalgebras of quantum mechanical observables. *International Journal of Theoretical Physics*, 38:2441–2484, 1999.
- [Halvorson, 2001] Hans Halvorson. On the nature of continuous physical quantities in classical and quantum mechanics. *Journal of Philosophical Logic*, 30:27–50, 2001.
- [Hardy, 1992] Lucien Hardy. Quantum mechanics, local realistic theories, and Lorentz-invariant realistic theories. *Physical Review Letters*, 68:2981–2984, 1992.

- [Hardy, 2001] Lucien Hardy. Quantum theory from five reasonable axioms. quant-ph/0101012, 2001.
- [Hardy, 2002] Lucien Hardy. Why quantum theory? In J. Butterfield and T. Placek, editors, *Proceedings of the NATO Advanced Research Workshop on Modality, Probability and Bell's Theorem*, Amsterdam, 2002. IOS Press.
- [Healey, 1989] Richard Healey. *The Philosophy of Quantum Mechanics: An Interactive Interpretation*. Cambridge University Press, Cambridge, 1989.
- [Heisenberg, 1927] Werner Heisenberg. Ueber den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Zeitschrift für Physik*, 43:172–198, 1927.
- [Heisenberg, 1930] Werner Heisenberg. *Die Physikalischen Prinzipien der Quantenmechanik*. Hirzel, Leipzig, 1930. English translation *The Physical Principles of Quantum Theory* published by University of Chicago Press, Chicago.
- [Hooker, 1975] C. A. Hooker, editor. *The Logico-Algebraic Approach to Quantum Mechanics*, volume 1. D. Reidel, Dordrecht, 1975.
- [Hooker, 1979] C. A. Hooker, editor. *The Logico-Algebraic Approach to Quantum Mechanics*, volume 2. D. Reidel, Dordrecht, 1979.
- [Howard, 2003] Don Howard. What makes a concept classical? <http://www.nd.edu/~dhoward1/Classcon.pdf>, 2003.
- [Jarrett, 1984] Jon Jarrett. On the physical significance of the locality conditions in the Bell arguments. *Noûs*, 18:569–589, 1984.
- [Jauch, 1968] J. M. Jauch. *Foundations of Quantum Mechanics*. Addison Wesley, Reading, MA, 1968.
- [Joos and Zeh, 1985] E. Joos and H. D. Zeh. The emergence of classical properties through interaction with the environment. *Zeitschrift für Physik*, B59:223–243, 1985.
- [Joos, 1986] E. Joos. Quantum theory and the appearance of a classical world. In D. Greenberger, editor, *New Techniques and Ideas in Quantum Measurement Theory*, volume 480 of *Annals of the New York Academy of Science*, pages 6–13, New York, 1986. New York Academy of Science.
- [Kochen and Specker, 1967] S. Kochen and E. P. Specker. The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17:59–87, 1967.
- [Kochen, 1985] S. Kochen. A new interpretation of quantum mechanics. In P. Mittelstaedt and P. Lahti, editors, *Symposium on the Foundations of Modern Physics*, pages 923–928, Singapore, 1985. World Scientific.
- [Landsman, 1998] N. P. Landsman. *Mathematical Topics Between Classical and Quantum Mechanics*. Springer-Verlag, New York, 1998.
- [Leavens, 1990] C. Leavens. Transmission, reflection, and dwell times within Bohm's causal interpretation of quantum mechanics. *Solid State Communications*, 74:923–928, 1990.
- [Lewis, 1973] David Lewis. Causation. *Journal of Philosophy*, 70:556–567, 1973.
- [Ludwig, 1983] G. Ludwig. *Mathematical Foundations of Quantum Mechanics*. Springer-Verlag, New York, 1983.
- [Ma et al., 1998] C. Ma, E. F. Arias, T. M. Eubanks, A. L. Fey, A-M. Gontier, C. S. Jacobs, O.J. Sovers, B. A. Archinal, and P. Charlot. The international celestial reference frame as realized by very long baseline interferometry. *The Astronomical Journal*, 116:516–546, 1998.
- [Mackey, 1949] G. W. Mackey. Imprimitivity for representations of locally compact groups. *Proceedings of the National Academy of Sciences*, 35:156–162, 1949.
- [Mackey, 1978] G. W. Mackey. *Unitary Group Representations in Physics Probability and Number Theory*, volume 55 of *Math Lecture Notes Series*. Benjamin Cummings Publishing Company, Menlo Park, CA, 1978.
- [Mackey, 1996] G. W. Mackey. The relationship between classical and quantum mechanics. In L. A. Coburn and M. A. Rieffel, editors, *Perspectives on Quantization*. American Mathematical Society, Providence, RI, 1996.
- [Massimi, 2004] Michela Massimi. *Pauli's Exclusion Principle: The Origin and Validation of a Scientific Principle*. Cambridge University Press, Cambridge, 2004.
- [Maudlin, 1994] Tim Maudlin. *Quantum Non-Locality and Relativity*. Blackwell, Oxford, 1994.
- [Mermin, 1990] David Mermin. Simple unified form for the major no-hidden variables theorems. *Physical Review Letters*, 65:3373–3376, 1990.
- [Mittelstaedt, 1998] Peter Mittelstaedt. *The Interpretation of Quantum Mechanics and the Measurement Process*. Cambridge University Press, Cambridge, 1998.

- [Pauli, 1958] W. Pauli. *Die allgemeinen Prinzipien der Wellenmechanik*, volume 1 of *Handbuch der Physik*, pages 1–168. Springer, Berlin, 1958.
- [Pearle, 1989] P. Pearle. Combining stochastic dynamical state-vector reduction with spontaneous localization. *Physical Review A*, 39:2277–2289, 1989.
- [Piron, 1964] C. Piron. Axiomatique quantique. *Helvetica Physica Acta*, 37:439–468, 1964.
- [Piron, 1976] C. Piron. *Foundations of Quantum Physics*. W.A. Benjamin, Reading, MA, 1976.
- [Pitowsky, 1991] Itamar Pitowsky. Bohm quantum potentials and quantum-gravity. *Foundations of Physics*, 21:343–352, 1991.
- [Primas, 1990] H. Primas. Induced nonlinear time evolution of open quantum objects. In A. Miller, editor, *Sixty-Two Years of Uncertainty*, pages 259–280. Plenum, New York, 1990.
- [Putnam, 1969] Hilary Putnam. Is logic empirical? In *M. Wartofsky and R. Cohen*, volume 5 of *Boston Studies in the Philosophy of Science*, pages 181–206. Humanities Press, New York, 1969.
- [Reichenbach, 1956] Hans Reichenbach. *The Direction of Time*. University of Los Angeles Press, Berkeley, 1956.
- [Rellich, 1946] F. Rellich. Der Eindeutigkeitssatz für die Lösungen der Quantenmechanischen Vertauschungsrelationen. *Nachrichten Akad. Wiss. Göttingen*, Math. Phys. Klasse II, 11A:107–115, 1946.
- [Riesz and Sz.-Nagy, 1955] Frigyes Riesz and Béla Sz.-Nagy. *Functional Analysis*. Frederick Ungar Publishing Co., New York, 1955.
- [Rudin, 1973] W. Rudin. *Functional Analysis*. McGraw-Hill, New York, 1973.
- [Sakurai, 1985] J. J. Sakurai. *Modern Quantum Mechanics*. Benjamin/Cummings, Menlo Park, CA, 1985.
- [Salmon, 1984] Wesley Salmon. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ, 1984.
- [Schiff, 1968] Leonard Schiff. *Quantum Mechanics*. McGraw-Hill, New York, 3rd edition, 1968.
- [Schrödinger, 1935a] Erwin Schrödinger. Die gegenwärtige Situation in der Quantenmechanik. *Naturwissenschaften*, 23:807–812; 823–828; 844–849, 1935. Translated as ‘The Present Situation in Quantum Mechanics’ by John D. Trimmer, *Proceedings of the American Philosophical Society*, 124, 323–338.
- [Schrödinger, 1935b] Erwin Schrödinger. Discussion of probability relations between separated systems. *Proceedings of the Cambridge Philosophical Society*, 31:555–563, 1935.
- [Scully and Walther, 1989] Marlan O. Scully and Herbert Walther. Quantum optical test of observation and complementarity in quantum mechanics. *Physical Review A*, 39:5229–5236, 1989.
- [Shimony, 1978] Abner Shimony. Metaphysical problems in the foundations of quantum mechanics. *International Philosophical Quarterly*, 18:3–17, 1978.
- [Shimony, 1986] Abner Shimony. Events and processes in the quantum world. In R. Penrose and C. Isham, editors, *Quantum Concepts in Space and Time*, pages 182–203. Clarendon Press, Oxford, 1986.
- [Solèr, 1995] M. P. Solèr. Characterization of Hilbert spaces with orthomodular spaces. *Communications in Algebra*, 23:219–243, 1995.
- [Stone, 1932] Marshall Stone. On one-parameter unitary groups in Hilbert space. *Annals of Mathematics*, 33:643–648, 1932.
- [Summers, 2001] Stephen Summers. On the Stone-von Neumann uniqueness theorem and its ramifications. In Miklos Rédei and Michael Stöltzner, editors, *John von Neumann and the Foundations of Quantum Physics*, volume 8 of *Institute Vienna Circle Yearbook*, pages 135–152. Dordrecht, 2001. Kluwer.
- [Szabó and Fine, 2002] László Szabó and Arthur Fine. A local hidden variable theory for the GHZ experiment. *Physics Letters A*, 295:229–240, 2002.
- [Tanona, 2002] Scott Tanona. *From Correspondence to Complementarity: The Emergence of Bohr’s Copenhagen Interpretation of Quantum Mechanics*. PhD thesis, Indiana University, 2002.
- [Tanona, 2006] Scott Tanona. Theory, coordination, and empirical meaning in modern physics. In Mary Domski and Michael Dickson, editors, *Synthesis and the Growth of Knowledge*. Open Court, Peoria, IL, 2006.
- [Uffink, 1994] Jos Uffink. Two new kinds of uncertainty relations. In D. Han, editor, *Proceedings of the Third international Workshop on Squeezed States and Uncertainty Relations*, volume 3270 of *NASA Conference Publications*, pages 155–161, Washington, DC, 1994. NASA.

- [Uffink, 1999a] Jos Uffink. How to protect the interpretation of the wave function against protective measurements. *Physical Review A*, 60:3474–3481, 1999.
- [Uffink, 1999b] Jos Uffink. The principle of the common cause faces the Bernstein paradox. *Philosophy of Science*, 66:S512–S525, 1999.
- [Valentini, 1991a] A. Valentini. Signal-locality, uncertainty and the subquantum J-theorem. I. *Physics Letters A*, 156:5–11, 1991.
- [Valentini, 1991b] A. Valentini. Signal-locality, uncertainty and the subquantum J-theorem. II. *Physics Letters A*, 158:1–8, 1991.
- [van Fraassen, 1972] Bas van Fraassen. A formal approach to the philosophy of science. In R. Colodny, editor, *Paradigms and Paradoxes: The Philosophical Challenge of the Quantum Domain*, pages 303–366, Pittsburgh, 1972. University of Pittsburgh Press.
- [Varadarajan, 1985] V. S. Varadarajan. *Geometry of Quantum Theory*. Springer-Verlag, Berlin, 1985.
- [Vermaas, 2000] Pieter Vermaas. *A Philosopher's Understanding of Quantum Mechanics: Possibilities and Impossibilities of a Modal Interpretation*. Cambridge University Press, Cambridge, 2000.
- [von Neumann, 1931] J. von Neumann. Die Eindeutigkeit der Schrödingerschen Operatoren. *Mathematische Annalen*, 104:570–578, 1931.
- [von Neumann, 1932] J. von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Springer, Berlin, 1932. English translation *Mathematical Foundations of Quantum Theory* published by Princeton University Press, Princeton, 1955.
- [Walborn *et al.*, 2002] S. P. Walborn, M. O. Terra Cunha, S. Pádua, and C. H. Monken. Double-slit quantum eraser. *Physical Review A*, 65:033818, 2002.
- [Wigner, 1931] Eugene P. Wigner. *Gruppentheorie und ihre Anwendung auf die Quantenmechanik der Atomspektren*. Vieweg, Braunschweig, 1931.
- [Wigner, 1961] Eugene P. Wigner. Remarks on the mind-body question. In I. J. Good, editor, *The Scientist Speculates*, pages 284–302, London, 1961. Heinemann.

This page intentionally left blank

BETWEEN CLASSICAL AND QUANTUM

N.P. Landsman

‘But the worst thing is that I am quite unable to clarify the transition [of matrix mechanics] to the classical theory.’ [Heisenberg to Pauli, October 23th, 1925]¹

‘Hendrik Lorentz considered the establishment of the correct relation between the classical and the quantum theory as the most fundamental problem of future research. This problem bothered him as much as it did Planck.’ [Mehra and Rechenberg, 2000, 721]

‘Thus quantum mechanics occupies a very unusual place among physical theories: it contains classical mechanics as a limiting case, yet at the same time it requires this limiting case for its own formulation.’ [Landau and Lifshitz, 1977, 3]

1 INTRODUCTION

Most modern physicists and philosophers would agree that a decent interpretation of quantum mechanics should fulfill at least two criteria. Firstly, it has to elucidate the physical meaning of its mathematical formalism and thereby secure the empirical content of the theory. This point (which we address only in a derivative way) was clearly recognized by all the founders of quantum theory.² Secondly (and this *is* the subject of this paper), it has to explain at least the *appearance* of the classical world.³ As shown by our second quotation above, Planck saw the difficulty this poses, and as a first contribution he noted that the high-temperature

¹‘Aber das Schlimmste ist, daß ich über den Übergang in die klassische Theorie nie Klarheit bekommen kann.’ See [Pauli, 1979, 251].

²The history of quantum theory has been described in a large number of books. The most detailed presentation is in [Mehra and Rechenberg, 1982–2001], but this multi-volume series has by no means superseded smaller works such as [Jammer, 1966; van der Waerden, 1967; Hendry, 1984; Darrigol, 1992], and [Beller, 1999]. Much information may also be found in biographies such as [Heisenberg, 1969; Pais, 1982; Moore, 1989; Pais, 1991; Cassidy, 1992; Heilbron, 2000; Enz, 2002], etc. See also [Pauli, 1979]. A new project on the history of matrix mechanics led by Jürgen Renn is on its way.

³That these points are quite distinct is shown by the Copenhagen Interpretation, which exclusively addresses the first at utter neglect of the second. Nonetheless, in most other approaches to quantum mechanics there is substantial overlap between the various mechanisms that are proposed to fulfill the two criteria in question.

limit of his formula for black-body radiation converged to the classical expression. Although Bohr believed that *quantum mechanics should be interpreted through classical physics*, among the founders of the theory he seems to have been unique in his lack of appreciation of the problem of *deriving classical physics from quantum theory*. Nonetheless, through his correspondence principle (which he proposed in order to address the *first* problem above rather than the second) Bohr made one of the most profound contributions to the issue. Heisenberg initially recognized the problem, but quite erroneously came to believe he had solved it in his renowned paper on the uncertainty relations.⁴ Einstein famously did not believe in the fundamental nature of quantum theory, whereas Schrödinger was well aware of the problem from the beginning, later highlighted the issue with his legendary cat, and at various stages in his career made important technical contributions towards its resolution. Ehrenfest stated the well-known theorem named after him. Von Neumann saw the difficulty, too, and addressed it by means of his well-known analysis of the measurement procedure in quantum mechanics.

The problem is actually even more acute than the founders of quantum theory foresaw. The experimental realization of Schrödinger's cat is nearer than most physicists would feel comfortable with [Leggett, 2002; Brezger *et al.*, 2002; Chiorescu *et al.*, 2003; Marshall *et al.*, 2003; Devoret *et al.*, 2004]. Moreover, awkward superpositions are by no means confined to physics laboratories: due to its chaotic motion, Saturn's moon Hyperion (which is about the size of New York) has been estimated to spread out all over its orbit within 20 years if treated as an isolated quantum-mechanical wave packet [Zurek and Paz, 1995]. Furthermore, decoherence theorists have made the point that "measurement" is not only a procedure carried out by experimental physicists in their labs, but takes place in Nature all the time without any human intervention. On the conceptual side, parties as diverse as Bohm and Bell and their followers on the one hand and the quantum cosmologists on the other have argued that a "Heisenberg cut" between object and observer cannot possibly lie at the basis of a fundamental theory of physics.⁵ These and other remarkable insights of the past few decades have drawn wide attention to the importance of the problem of interpreting quantum mechanics, and in particular of explaining classical physics from it.

We will discuss these ideas in more detail below, and indeed our discussion of the relationship between classical and quantum mechanics will be partly historical. However, other than that it will be technical and mathematically rigorous. For the problem at hand is so delicate that in this area sloppy mathematics is almost guaranteed to lead to unreliable physics and conceptual confusion (notwithstand-

⁴'One can see that the transition from micro- to macro-mechanics is now very easy to understand: classical mechanics is altogether part of quantum mechanics.' (Heisenberg to Bohr, 19 March 1927, just before the submission on 23 March of Heisenberg [1927]. See *Bohr's Scientific Correspondence in the Archives for the History of Quantum Physics*).

⁵Not to speak of the problem, also raised by quantum cosmologists, of deriving classical space-time from some theory of quantum gravity. This is certainly part of the general program of deriving classical physics from quantum theory, but unfortunately it cannot be discussed in this paper.

ing the undeniable success of poor man's math elsewhere in theoretical physics). Except for von Neumann, this was not the attitude of the pioneers of quantum mechanics; but while it has to be acknowledged that many of their ideas are still central to the current discussion, these ideas *per se* have *not* solved the problem. Thus we assume the reader to be familiar with the Hilbert space formalism of quantum mechanics,⁶ and for some parts of this paper (notably Section 6 and parts of Section 4) also with the basic theory of C^* -algebras and its applications to quantum theory.⁷ In addition, some previous encounter with the conceptual problems of quantum theory would be helpful.⁸

Which ideas *have* solved the problem of explaining the appearance of the classical world from quantum theory? In our opinion, none have, although since the founding days of quantum mechanics a number of new ideas have been proposed that almost certainly will play a role in the eventual resolution, should it ever be found. These ideas surely include:

- The limit $\hbar \rightarrow 0$ of small Planck's constant (coming of age with the mathematical field of microlocal analysis);
- The limit $N \rightarrow \infty$ of a large system with N degrees of freedom (studied in a serious way only after the emergence of C^* -algebraic methods);
- Decoherence and consistent histories.

Mathematically, the second limit may be seen as a special case of the first, though the underlying physical situation is of course quite different. In any case, after a detailed analysis our conclusion will be that none of these ideas in isolation is capable of explaining the classical world, but that there is some hope that by combining all three of them, one might do so in the future.

Because of the fact that the subject matter of this review is unfinished business, to date one may adopt a number of internally consistent but mutually incompatible philosophical stances on the relationship between classical and quantum theory. Two extreme ones, which are always useful to keep in mind whether one holds one of them or not, are:

1. Quantum theory is fundamental and universally valid, and the classical world has only "relative" or "perspectival" existence.
2. Quantum theory is an approximate and derived theory, possibly false, and the classical world exists absolutely.

⁶Apart from seasoned classics such as [Mackey, 1963; Jauch, 1968; Prugovecki, 1971; Reed and Simon, 1972], or [Thirring, 1981], the reader might consult more recent books such as [Gustafson and Sigal, 2003] or [Williams, 2003]. See also [Dickson, 2005].

⁷For physics-oriented introductions to C^* -algebras see [Davies, 1976; Roberts and Roepstorff, 1969; Primas, 1983; Thirring, 1983; Emch, 1984; Strocchi, 1985; Sewell, 1986; Roberts, 1990; Haag, 1992; Landsman, 1998; Araki, 1999], and [Sewell, 2002]. Authoritative mathematical texts include [Kadison and Ringrose, 1983; 1986] and [Takesaki, 2003].

⁸Trustworthy books include, for example, [Scheibe, 1973; Jammer, 1974; van Fraassen, 1991; d'Espagnat, 1995; Peres, 1995; Omnès, 1994; 1999; Bub, 1997], and [Mittelstaedt, 2004].

An example of a position that our modern understanding of the measurement problem⁹ has rendered internally inconsistent is:

3. Quantum theory is fundamental and universally valid, and (yet) the classical world exists absolutely.

In some sense stance 1 originates with Heisenberg [1927], but the modern era started with Everett [1957].¹⁰ These days, most decoherence theorists, consistent historians, and modal interpreters seem to support it. Stance 2 has a long and respectable pedigree unequivocally, including among others Einstein, Schrödinger, and Bell. More recent backing has come from Leggett as well as from “spontaneous collapse” theorists such as Pearle, Ghirardi, Rimini, Weber, and others. As we shall see in Section 3, Bohr’s position eludes classification according to these terms; our three stances being of an ontological nature, he probably would have found each of them unattractive.

Of course, one has to specify what the terminology involved means. By quantum theory we mean standard quantum mechanics including the eigenvector-eigenvalue link.¹¹ Modal interpretations of quantum mechanics [Dieks, 1989a; 1989b; van Fraassen, 1991; Bub, 1999; Vermaas, 2000; Bene & Dieks, 2002; Dickson, 2005] deny this link, and lead to positions close to or identical to stance 1. The projection postulate is neither endorsed nor denied when we generically speak of quantum theory.

It is a bit harder to say what “the classical world” means. In the present discussion we evidently can *not* define the classical world as the world that exists independently of observation – as Bohr did, see Subsection 3.1 – but neither can it be taken to mean the part of the world that is described by the laws of classical physics full stop; for if stance 1 is correct, then these laws are only approximately valid, if at all. Thus we simply put it like this:

The classical world is what observation shows us to behave – with appropriate accuracy – according to the laws of classical physics.

There should be little room for doubt as to what ‘with appropriate accuracy’ means: the existence of the colour grey does not imply the nonexistence of black and white!

We *can* define the *absolute existence* of the classical world à la Bohr as its existence independently of observers or measuring devices. Compare with Moore’s [1939] proof of the existence of the external world:

⁹See the books cited in footnote 8, especially [Mittelstaedt, 2004].

¹⁰Note, though, that stance 1 by no means implies the so-called Many-Worlds Interpretation, which also in our opinion is ‘simply a meaningless collage of words’ [Leggett, 2002].

¹¹Let A be a selfadjoint operator on a Hilbert space \mathcal{H} , with associated projection-valued measure $P(\Delta)$, $\Delta \subset \mathbb{R}$, so that $A = \int dP(\lambda) \lambda$ (see also footnote 99 below). The eigenvector-eigenvalue link states that a state Ψ of the system lies in $P(\Delta)\mathcal{H}$ if and only if A takes some value in Δ for sure. In particular, if Ψ is an eigenvector of A with eigenvalue λ (so that $P(\{\lambda\}) \neq 0$ and $\Psi \in P(\{\lambda\})\mathcal{H}$), then A takes the value λ in the state Ψ with probability one. In general, the probability $p_\Psi(\Delta)$ that in a state Ψ the observable a takes some value in Δ (“upon measurement”) is given by the Born–von Neumann rule $p_\Psi(\Delta) = (\Psi, P(\Delta)\Psi)$.

How? By holding up my two hands, and saying, as I make a certain gesture with the right hand, ‘Here is one hand’, and adding, as I make a certain gesture with the left, ‘and here is another’.

Those holding position 1, then, maintain that *the classical world exists only as an appearance relative to a certain specification*, where the specification in question could be an observer (Heisenberg), a certain class of observers and states (as in decoherence theory), or some coarse-graining of the Universe defined by a particular consistent set of histories, etc. If the notion of an observer is construed in a sufficiently abstract and general sense, one might also formulate stance 1 as claiming that the classical world merely exists *from the perspective of the observer* (or the corresponding class of observables).¹² For example, Schrödinger’s cat “paradox” dissolves at once when the appropriate perspective is introduced; cf. Subsection 6.6.

Those holding stance 2, on the other hand, believe that the classical world exists in an absolute sense (as Moore did). Thus stance 2 is akin to common-sense realism, though the distinction between 1 and 2 is largely independent of the issue of scientific realism.¹³ For defendants of stance 1 usually still believe in the existence of some observer-independent reality (namely somewhere in the quantum realm), but deny that this reality incorporates the world observed around us. This justifies a pretty vague specification of such an important notion as the classical world: one of the interesting outcomes of the otherwise futile discussions surrounding the Many Worlds Interpretation has been the insight that *if quantum mechanics is fundamental, then the notion of a classical world is intrinsically vague and approximate*. Hence it would be self-defeating to be too precise at this point.¹⁴

Although stance 1 is considered defensive if not cowardly by adherents of stance 2, it is a highly nontrivial mathematical fact that so far it seems supported by the formalism of quantum mechanics. In his derision of what he called ‘FAPP’ (= For All Practical Purposes) solutions to the measurement problem (and more general attempts to explain the appearance of the classical world from quantum theory), Bell [1987; 2001] and others in his wake mistook a profound epistemological stance for a poor defensive move.¹⁵ It is, in fact, stance 2 that we would recommend to the cowardly: for proving or disproving stance 1 seems the real challenge of the

¹²The terminology “perspectival” was suggested to the author by Richard Healey.

¹³See [Landsman, 1995] for a more elaborate discussion of realism in this context. Words like “objective” or “subjective” are not likely to be helpful in drawing the distinction either: the claim that ‘my children are the loveliest creatures in the world’ is at first glance subjective, but it can trivially be turned into an objective one through the reformulation that ‘Klaas Landsman finds his children the loveliest creatures in the world’. Similarly, the proposition that (perhaps due to decoherence) ‘local observers find that the world is classical’ is perfectly objective, although it describes a subjective experience. See also [Davidson, 2001].

¹⁴See [Wallace, 2002; 2003]; also cf. [Butterfield, 2002]. This point was not lost on Bohr and Heisenberg either; see [Scheibe, 1973].

¹⁵The insistence on “precision” in such literature is reminiscent of Planck’s long-held belief in the absolute nature of irreversibility [Darrigol, 1992; Heilbron, 2002]. It should be mentioned that although Planck’s stubbornness by historical accident led him to take the first steps towards quantum theory, he eventually gave it up to side with Boltzmann.

entire debate, and we regard the technical content of this paper as a survey of progress towards actually proving it. Indeed, to sum up our conclusions, we claim that there is good evidence that:

1. Classical physics emerges from quantum theory in the limit $\hbar \rightarrow 0$ or $N \rightarrow \infty$ *provided that the system is in certain “classical” states and is monitored with “classical” observables only*;
2. Decoherence and consistent histories will probably explain *why* the system happens to be in such states and has to be observed in such a way.

However, even if one fine day this scheme will be made to work, the explanation of the appearance of the classical world from quantum theory will be predicated on an external solution of the notorious ‘from “and” to “or” problem’: If quantum mechanics predicts various possible outcomes with certain probabilities, why does only *one* of these appear to us?¹⁶

For a more detailed outline of this paper we refer to the table of contents above. Most philosophical discussion will be found in Section 3 on the Copenhagen interpretation, since whatever its merits, it has undeniably set the stage for the entire discussion on the relationship between classical and quantum.¹⁷ The remainder of the paper will be of an almost purely technical nature. Beyond this point we will try to avoid controversy, but when unavoidable it will be confined to the Epilogues appended to Sections 3-6. The final Epilogue (Section 8) expresses our deepest thoughts on the subject.

2 EARLY HISTORY

This section is a recapitulation of the opinions and contributions of the founders of quantum mechanics regarding the relationship between classical and quantum. More detail may be found in the books cited in footnote 2 and in specific literature to be cited; for an impressive (but incomplete) bibliography see also [Gutzwiller, 1998]. The early history of quantum theory is of interest in its own right, concerned as it is with one of the most significant scientific revolutions in history. Although this history is not a main focus of this paper, it is of special significance for our theme. For the usual and mistaken interpretation of Planck’s work (i.e. the idea that he introduced something like a “quantum postulate”, see Subsection 3.2 below) appears to have triggered the belief that quantum theory and Planck’s constant are related to a universal discontinuity in Nature. Indeed, this discontinuity is sometimes even felt to mark the basic difference between classical and quantum physics. This belief is particularly evident in the writings of Bohr, but still resonates even today.

¹⁶It has to be acknowledged that we owe the insistence on this question to the defendants of stance 2. See also footnote 10.

¹⁷We do not discuss the classical limit of quantum mechanics in the philosophical setting of theory reduction and intertheoretic relations; see, e.g., [Scheibe, 1999] and [Batterman, 2002].

2.1 Planck and Einstein

The relationship between classical physics and quantum theory is so subtle and confusing that historians and physicists cannot even agree about the precise way the classical gave way to the quantum! As Darrigol [2001] puts it: ‘During the past twenty years, historians [and physicists] have disagreed over the meaning of the quanta which Max Planck introduced in his black-body theory of 1900. The source of this confusion is the publication (...) of Thomas Kuhn’s [1978] iconoclastic thesis that Planck did not mean his energy quanta to express a quantum discontinuity.’

As is well known (cf. [Mehra and Rechenberg, 1982a], etc.), Planck initially derived Wien’s law for blackbody radiation in the context of his (i.e. Planck’s) program of establishing the absolute nature of irreversibility (competing with Boltzmann’s probabilistic approach, which eventually triumphed). When new high-precision measurements in October 1900 turned out to refute Wien’s law, Planck first guessed his expression

$$(1) \quad E_\nu/N_\nu = h\nu/(e^{h\nu/kT} - 1)$$

for the correct law, *en passant* introducing two new constants of nature h and k ,¹⁸ and subsequently, on December 14, 1900, presented a theoretical derivation of his law in which he allegedly introduced the idea that the energy of the resonators making up his black body was quantized in units of $\varepsilon_\nu = h\nu$ (where ν is the frequency of a given resonator). This derivation is generally seen as the birth of quantum theory, with the associated date of birth just mentioned.

However, it is clear by now [Kuhn, 1978; Darrigol, 1992; 2001; Carson, 2000; Brush, 2002] that Planck was at best agnostic about the energy of his resonators, and at worst assigned them a continuous energy spectrum. Technically, in the particular derivation of his empirical law that eventually turned out to lead to the desired result (which relied on Boltzmann’s concept of entropy),¹⁹ Planck had to count the number of ways a given amount of energy E_ν could be distributed over a given number of resonators N_ν at frequency ν . This number is, of course, infinite, hence in order to find a finite answer Planck followed Boltzmann in breaking up E_ν into a large number A_ν of portions of identical size ε_ν , so that $A_\nu\varepsilon_\nu = E_\nu$.²⁰ Now, as we all know, whereas Boltzmann let $\varepsilon_\nu \rightarrow 0$ at the end of his corresponding calculation for a gas, Planck discovered that his empirical blackbody law emerged if he assumed the relation $\varepsilon_\nu = h\nu$.

However, this postulate did *not* imply that Planck quantized the energy of his resonators. In fact, in his definition of a given distribution he counted the number of resonators with energy *between* say $(k - 1)\varepsilon_\nu$ and $k\varepsilon_\nu$ (for some $k \in$

¹⁸Hence Boltzmann’s constant k was introduced by Planck, who was the first to write down the formula $S = k \log W$.

¹⁹Despite the fact that Planck only converted to Boltzmann’s approach to irreversibility around 1914.

²⁰The number in question is then given by $(N + A - 1)/(N - 1)!A!$, dropping the dependence on ν in the notation.

\mathbb{N}), as Boltzmann did in an analogous way for a gas, rather than the number of resonators with energy $k\varepsilon_\nu$, as most physicists came to interpret his procedure. More generally, there is overwhelming textual evidence that Planck himself by no means believed or implied that he had quantized energy; for one thing, in his Nobel Prize Lecture in 1920 he attributed the correct interpretation of the energy-quanta ε_ν to Einstein. Indeed, the modern understanding of the earliest phase of quantum theory is that it was Einstein rather than Planck who, during the period 1900–1905, clearly realized that Planck’s radiation law marked a break with classical physics [Büttner *et al.*, 2003]. This insight, then, led Einstein to the quantization of energy. This he did in a twofold way, both in connection with Planck’s resonators — interpreted by Einstein as harmonic oscillators in the modern way — and, in a closely related move, through his concept of a photon. Although Planck of course introduced the constant named after him, and as such is the founding *father* of the theory characterized by \hbar , it is the introduction of the photon that made Einstein at least the *mother* of quantum theory. Einstein himself may well have regarded the photon as his most revolutionary discovery, for what he wrote about his pertinent paper is not matched in self-confidence by anything he said about relativity: ‘Sie handelt über die Strahlung und die energetischen Eigenschaften des Lichtes und ist sehr revolutionär.’²¹

Finally, in the light of the present paper, it deserves to be mentioned that Einstein [1905] and Planck [1906] were the first to comment on the classical limit of quantum theory; see the preamble to Section 5 below.

2.2 Bohr

Bohr’s brilliant model of the atom reinforced his idea that quantum theory was a theory of quanta.²² Since this model simultaneously highlighted the clash between classical and quantum physics *and* carried the germ of a resolution of this conflict through Bohr’s equally brilliant correspondence principle, it is worth saying a few words about it here.²³ Bohr’s atomic model addressed the radiative instability of Rutherford’s solar-system-style atom:²⁴ according to the electrodynamics of Lorentz, an accelerating electron should radiate, and since the envisaged circular or elliptical motion of an electron around the nucleus is a special case of an accelerated motion, the electron should continuously lose energy and spiral towards

²¹‘[This paper] is about radiation and the energetic properties of light, and is very revolutionary.’ See also the Preface to Pais [1982].

²²Although at the time Bohr followed practically all physicists in their rejection of Einstein’s photon, since he believed that during a quantum jump the atom emits electromagnetic radiation in the form of a spherical wave. His model probably would have gained in consistency by adopting the photon picture of radiation, but in fact Bohr was to be the last prominent opponent of the photon, resisting the idea until 1925. See also Blair Bolles [2004] and footnote 33 below.

²³Cf. [Darrigol, 1992] for a detailed treatment; also see [Liboff, 1984] and [Steiner, 1998].

²⁴The solar system provides the popular visualization of Rutherford’s atom, but his own picture was more akin to Saturn’ rings than to a planet orbiting the Sun.

the nucleus.²⁵ Bohr countered this instability by three simultaneous moves, each of striking originality:

1. He introduced a quantization condition that singled out only a discrete number of allowed electronic orbits (which subsequently were to be described using classical mechanics, for example, in Bohr's calculation of the Rydberg constant R).
2. He replaced the emission of continuous radiation called for by Lorentz by quantum jumps with unpredictable destinations taking place at unpredictable moments, during which the atom emits light with energy equal to the energy difference of the orbits between which the electron jumps.
3. He prevented the collapse of the atom through such quantum jumps by introducing the notion of a ground state, below which no electron could fall.

With these postulates, for which at the time there existed no foundation whatsoever,²⁶ Bohr explained the spectrum of the hydrogen atom, including an amazingly accurate calculation of R . Moreover, he proposed what was destined to be the key guiding principle in the search for quantum mechanics in the coming decade, viz. the correspondence principle (cf. [Darrigol, 1992, *passim*], and [Mehra and Rechenberg, 1982a, 249–257]).

In general, there is no relation between the energy that an electron loses during a particular quantum jump and the energy it would have radiated classically (i.e. according to Lorentz) in the orbit it revolves around preceding this jump. Indeed, in the ground state it cannot radiate through quantum jumps at all, whereas according to classical electrodynamics it should radiate all the time. However, Bohr saw that in the opposite case of very wide orbits (i.e. those having very large principal quantum numbers n), the frequency $\nu = (E_n - E_{n-1})/h$ (with $E_n = -R/n^2$) of the emitted radiation approximately corresponds to the frequency of the lowest harmonic of the classical theory, applied to electron motion in the initial orbit.²⁷ Moreover, the measured intensity of the associated spectral line (which theoretically should be related to the probability of the quantum jump, a quantity out of the reach of early quantum theory), similarly turned out to be given by classical electrodynamics. This property, which in simple cases could be verified either by explicit computation or by experiment, became a guiding principle in situations where it could not be verified, and was sometimes even

²⁵In addition, any Rutherford style atom with more than one electron is mechanically unstable, since the electrons repel each other, as opposed to planets, which attract each other.

²⁶What has hitherto been mathematically proved of Bohr's atomic model is the existence of a ground state (see [Griesemer *et al.*, 2001], and references therein for the greatest generality available to date) and the metastability of the excited states of the atom after coupling to the electromagnetic field (cf. [Bach *et al.*, 1998; 1999] and [Gustafson and Sigal, 2003]). The energy spectrum is discrete only if the radiation field is decoupled, leading to the usual computation of the spectrum of the hydrogen atom first performed by Schrödinger and Weyl. See also the end of Subsection 5.4.

²⁷Similarly, higher harmonics correspond to quantum jumps $n \rightarrow n - k$ for $k > 1$.

extended to low quantum numbers, especially when the classical theory predicted selection rules.

It should be emphasized that *Bohr's correspondence principle was concerned with the properties of radiation, rather than with the mechanical orbits themselves*.²⁸ This is not quite the same as what is usually called the correspondence principle in the modern literature.²⁹ In fact, although also this modern correspondence principle has a certain range of validity (as we shall see in detail in Section 5), Bohr never endorsed anything like that, and is even on record as opposing such a principle:³⁰

‘The place was Purcell’s office where Purcell and others had taken Bohr for a few minutes of rest [during a visit to the Physics Department at Harvard University in 1961]. They were in the midst of a general discussion when Bohr commented: “People say that classical mechanics is the limit of quantum mechanics when h goes to zero.” Then, Purcell recalled, Bohr shook his finger and walked to the blackboard on which he wrote e^2/hc . As he made three strokes under h , Bohr turned around and said, “you see h is in the denominator.”’

2.3 Heisenberg

Heisenberg’s [1925] paper *Über die quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen*³¹ is generally seen as a turning point in the development of quantum mechanics. Even A. Pais, no friend of Heisenberg’s,³² conceded that Heisenberg’s paper marked ‘one of the great jumps — perhaps the greatest — in the development of twentieth century physics.’ What did Heisenberg actually accomplish? This question is particularly interesting from the perspective of our theme.

At the time, atomic physics was in a state of crisis, to which various camps responded in different ways. Bohr’s approach might best be described as *damage control*: his quantum theory was a hybrid of classical mechanics adjusted by means of *ad hoc* quantization rules, whilst keeping electrodynamics classical at all cost.³³

²⁸As such, it remains to be verified in a rigorous way.

²⁹A typical example of the modern version is: ‘Non-relativistic quantum mechanics was founded on the correspondence principle of Bohr: “When the Planck constant \hbar can be considered small with respect to the other parameters such as masses and distances, quantum theory approaches classical Newton theory.”’ [Robert, 1998, 44]. The reference to Bohr is historically inaccurate!

³⁰Quoted from [Miller, 1984, 313].

³¹*On the quantum theoretical reinterpretation of kinematical and mechanical relations*. English translation in [van der Waerden, 1967].

³²For example, in [Pais, 2000], claiming to portray the ‘genius of science’, Heisenberg is conspicuously absent.

³³Continuing footnote 22, we quote from [Mehra and Rechenberg, 1982a, 256–257]: ‘Thus, in the early 1920s, Niels Bohr arrived at a definite point of view how to proceed forward in atomic theory. He wanted to make maximum use of what he called the “more dualistic prescription” (. . .) In it the atom was regarded as a mechanical system having discrete states and emitting radiation of discrete frequencies, determined (in a nonclassical way) by the energy differences

Einstein, who had been the first physicist to recognize the need to quantize classical electrodynamics, in the light of his triumph with General Relativity nonetheless dreamt of a classical field theory with singular solutions as the ultimate explanation of quantum phenomena. Born led the radical camp, which included Pauli: he saw the need for an entirely new mechanics replacing classical mechanics,³⁴ which was to be based on discrete quantities satisfying difference equations.³⁵ This was a leap in the dark, especially because of Pauli's frowning upon the correspondence principle [Hendry, 1984; Beller, 1999].

It was Heisenberg's genius to *interpolate* between Bohr and Born.³⁶ The meaning of his *Umdeutung* was to keep the classical equations of motion,³⁷ whilst reinterpreting the mathematical symbols occurring therein as (what were later recognized to be) matrices. Thus his *Umdeutung* $x \mapsto a(n, m)$ was a precursor of what now would be called a quantization map $f \mapsto Q_{\hbar}(f)$, where f is a classical observable, i.e. a function on phase space, and $Q_{\hbar}(f)$ is a quantum mechanical observable, in the sense of an operator on a Hilbert space or, more abstractly, an element of some C^* -algebra. See Section 4 below. As Heisenberg recognized, this move implies the noncommutativity of the quantum mechanical observables; it is this, rather than something like a "quantum postulate" (see Subsection 3.2 below), that is the defining characteristic of quantum mechanics. Indeed, most later work on quantum physics and practically all considerations on the connection between classical and quantum rely on Heisenberg's idea of *Umdeutung*. This even applies to the mathematical formalism as a whole; see Subsection 2.5.

We here use the term "observable" in a loose way. It is now well recognized [Mehra and Rechenberg, 1982b; Beller, 1999; Camilleri, 2005] that Heisenberg's claim that his formalism could be physically interpreted as the replacement of atomic orbits by observable quantities was a red herring, inspired by his discussions with Pauli. In fact, in quantum mechanics *any* mechanical quantity has to be "reinterpreted", whether or not it is observable. As Heisenberg [1969] recalls, Einstein reprimanded him for the illusion that physics admits an *a priori* notion of an observable, and explained that a theory determines what can be observed. Rethinking the issue of observability then led Heisenberg to his second major contribution to quantum mechanics, namely his uncertainty relations.

These relations were Heisenberg's own answer to the quote opening this paper. Indeed, matrix mechanics was initially an extremely abstract and formal scheme,

between stationary states; radiation, on the other hand, had to be described by the classical electrodynamic theory.'

³⁴It was Born who coined the name *quantum mechanics* even before Heisenberg's paper.

³⁵This idea had earlier occurred to Kramers.

³⁶Also literally! Heisenberg's traveled between Copenhagen and Göttingen most of the time.

³⁷This crucial aspect of *Umdeutung* was appreciated at once by Dirac [1926]: 'In a recent paper Heisenberg puts forward a new theory which suggests that it is not the equations of classical mechanics that are in any way at fault, but that the mathematical operations by which physical results are deduced from them require modification. (...) The correspondence between the quantum and classical theories lies not so much in the limiting agreement when $\hbar \rightarrow 0$ as in the fact that the mathematical operations on the two theories obey in many cases the same laws.'

which lacked not only any visualization but also the concept of a state (see below). Although these features were initially quite to the liking of Born, Heisenberg, Pauli, and Jordan, the success of Schrödinger's work forced them to renege on their radical stance, and look for a semiclassical picture supporting their mathematics; this was a considerable U-turn [Beller, 1999; Camilleri, 2005]. Heisenberg [1927] found such a picture, claiming that his uncertainty relations provided the 'intuitive content of the quantum theoretical kinematics and mechanics' (as his paper was called). His idea was that *the classical world emerged from quantum mechanics through observation*: 'The trajectory only comes into existence because we observe it.'³⁸ This idea was to become extremely influential, and could be regarded as the origin of stance 1 in the Introduction.

2.4 Schrödinger

The history of quantum mechanics is considerably clarified by the insight that Heisenberg and Schrödinger did not, as is generally believed, discover two equivalent formulations of the theory, but rather that Heisenberg [1925] identified the mathematical nature of the observables, whereas Schrödinger [1926a] found the description of states.³⁹ Matrix mechanics lacked the notion of a state, but by the same token wave mechanics initially had no observables; it was only in his attempts to relate wave mechanics to matrix mechanics that Schrödinger [1926c] introduced the position and momentum operators⁴⁰

$$(2) \quad \begin{aligned} \mathcal{Q}_{\hbar}(q^j) &= x^j; \\ \mathcal{Q}_{\hbar}(p_j) &= -i\hbar \frac{\partial}{\partial x^j}. \end{aligned}$$

This provided a new basis for Schrödinger's equation⁴¹

$$(3) \quad \left(-\frac{\hbar^2}{2m} \sum_{j=1}^n \frac{\partial^2}{\partial x_j^2} + V(x) \right) \Psi = i\hbar \frac{\partial \Psi}{\partial t},$$

by interpreting the left-hand side as $H\Psi$, with $H = \mathcal{Q}_{\hbar}(h)$ in terms of the classical Hamiltonian $h(p, q) = \sum_j p_j^2/2m + V(q)$. Thus Schrödinger founded the theory of the operators now named after him,⁴² and in doing so gave what is still the most important example of Heisenberg's idea of *Umdeutung* of classical observables.

³⁸'Die Bahn entsteht erst dadurch, daß wir sie beobachten.'

³⁹See also [Muller, 1997].

⁴⁰Here $j = 1, 2, 3$. In modern terms, the expressions on the right-hand side are unbounded operators on the Hilbert space $\mathcal{H} = L^2(\mathbb{R}^n)$. See Section 4 for more details. The expression x^j is a multiplication operator, i.e. $(x^j\Psi)(x) = x^j\Psi(x)$, whereas, obviously, $(\partial/\partial x^j\Psi)(x) = (\partial\Psi/\partial x^j)(x)$.

⁴¹Or the corresponding time-independent one, with $E\Psi$ on the right-hand side.

⁴²See [Reed and Simon, 1972; 1975; 1987; 1979; Cycon *et al.*, 1987; Hislop and Sigal, 1996; Hunziker and Sigal, 2000; Simon, 2000; Gustafson and Sigal, 2003]. For the mathematical origin of the Schrödinger equation also cf. [Simon, 1976].

Subsequently, correcting and expanding on certain ideas of Dirac, Pauli, and Schrödinger, von Neumann [1932] brilliantly glued these two parts together through the concept of a Hilbert space. He also gave an abstract form of the formulae of Born, Pauli, Dirac, and Jordan for the transition probabilities, thus completing the mathematical formulation of quantum mechanics.

However, this is not how Schrödinger saw his contribution. He intended wave mechanics as a full-fledged classical field theory of reality, rather than merely as one half (namely in modern parlance the state space half) of a probabilistic description of the world that still incorporated the quantum jumps he so detested [Mehra and Rechenberg, 1987; Götsch, 1992; Bitbol and Darrigol, 1992; Bitbol, 1996; Beller, 1999]. Particles were supposed to emerge in the form of wave packets, but it was immediately pointed out by Heisenberg, Lorentz, and others that in realistic situations such wave packets tend to spread in the course of time. This had initially been overlooked by Schrödinger [1926b], who had based his intuition on the special case of the harmonic oscillator. On the positive side, in the course of his unsuccessful attempts to derive classical particle mechanics from wave mechanics through the use of wave packets, Schrödinger [1926b] gave the first example of what is now called a *coherent state*. Here a quantum wave function Ψ_z is labeled by a ‘classical’ parameter z , in such a way that the quantum-mechanical time-evolution $\Psi_z(t)$ is approximately given by $\Psi_{z(t)}$, where $z(t)$ stands for some associated classical time-evolution; see Subsections 4.2 and 5.2 below. This has turned out to be a very important idea in understanding the transition from quantum to classical mechanics.

Furthermore, in the same paper Schrödinger [1926b] proposed the following wave-mechanical version of Bohr’s correspondence principle: classical atomic states should come from superpositions of a very large number (say at least 10,000) of highly excited states (i.e. energy eigenfunctions with very large quantum numbers). After decades of limited theoretical interest in this idea, interest in wave packets in atomic physics was revived in the late 1980s due to the development of modern experimental techniques based on lasers (such as pump-probing and phase-modulation). See [Robinett, 2004] for a recent technical review, or [Nauenberg *et al.*, 1994] for an earlier popular account. Roughly speaking, the picture that has emerged is this: a localized wave packet of the said type initially follows a time-evolution with almost classical periodicity, as Schrödinger hoped, but subsequently spreads out after a number of orbits. Consequently, during this second phase the probability distribution approximately fills the classical orbit (though not uniformly). Even more surprisingly, on a much longer time scale there is a phenomenon of *wave packet revival*, in which the wave packet recovers its initial localization. Then the whole cycle starts once again, so that one does see periodic behaviour, but not of the expected classical type. Hence even in what naively would be thought of as the thoroughly classical regime, wave phenomena continue to play a role, leading to quite unusual and unexpected behaviour. Although a rigorous mathematical description of wave packet revival has not yet been forthcoming, the overall picture (based on both “theoretical physics” style mathematics

and experiments) is clear enough.

It is debatable (and irrelevant) whether the story of wave packets has evolved according to Schrödinger's intentions (cf. [Littlejohn, 1986]); what is certain is that his other main idea on the relationship between classical and quantum has been extremely influential. This was, of course, Schrödinger's [1926a] "derivation" of his wave equation from the Hamilton–Jacobi formalism of classical mechanics. This gave rise to the WKB approximation and related methods; see Subsection 5.5.

In any case, where Schrödinger hoped for a classical interpretation of his wave function, and Heisenberg wanted to have nothing to do with it whatsoever [Beller, 1999], Born and Pauli were quick to realize its correct, probabilistic significance. Thus they deprived the wave function of its naive physical nature, and effectively degraded it to the purely mathematical status of a probability amplitude. And in doing so, Born and Pauli rendered the connection between quantum mechanics and classical mechanics almost incomprehensible once again! It was this incomprehensibility that Heisenberg addressed with his uncertainty relations.

2.5 von Neumann

Through its creation of the Hilbert space formalism of quantum mechanics, von Neumann's book [1932] can be seen as a mathematical implementation of Heisenberg's idea of *Umdeutung*. Von Neumann in effect proposed the following quantum-theoretical reinterpretations:

Phase space $M \mapsto$ Hilbert space \mathcal{H} ;

Classical observable (i.e. real-valued measurable function on M) \mapsto self-adjoint operator on \mathcal{H} ;

Pure state (seen as point in M) \mapsto unit vector (actually ray) in \mathcal{H} ;

Mixed state (i.e. probability measure on M) \mapsto density matrix on \mathcal{H} ;

Measurable subset of $M \mapsto$ closed linear subspace of \mathcal{H} ;

Set complement \mapsto orthogonal complement;

Union of subsets \mapsto closed linear span of subspaces;

Intersection of subsets \mapsto intersection of subspaces;

Yes-no question (i.e. characteristic function on M) \mapsto projection operator.⁴³

Here we assume for simplicity that quantum observables R on a Hilbert space \mathcal{H} are bounded operators, i.e. $R \in \mathcal{B}(\mathcal{H})$. Von Neumann actually *derived* his *Umdeutung* of classical mixed states as density matrices from his axiomatic characterization of quantum-mechanical states as linear maps $\text{Exp} : \mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C}$ that

⁴³Later on, he of course added the *Umdeutung* of a Boolean lattice by a modular lattice, and the ensuing *Umdeutung* of classical logic by quantum logic [Birkhoff and von Neumann, 1936].

satisfy $\text{Exp}(R) \geq 0$ when $R \geq 0$,⁴⁴ $\text{Exp}(1) = 1$,⁴⁵ and countable additivity on a commuting set of operators. For he proved that such a map Exp is necessarily given by a density matrix ρ according to $\text{Exp}(R) = \text{Tr}(\rho R)$.⁴⁶ A unit vector $\Psi \in \mathcal{H}$ defines a pure state in the sense of von Neumann, which we call ψ , by $\psi(R) = (\Psi, R\Psi)$ for $R \in \mathcal{B}(\mathcal{H})$. Similarly, a density matrix ρ on \mathcal{H} defines a (generally mixed) state, called ρ as well, by $\rho(R) = \text{Tr}(\rho R)$. In modern terminology, a state on $\mathcal{B}(\mathcal{H})$ as defined by von Neumann would be called a *normal* state. In the C^* -algebraic formulation of quantum physics (cf. footnote 7), this axiomatization has been maintained until the present day; here $\mathcal{B}(\mathcal{H})$ is replaced by more general algebras of observables in order to accommodate possible superselection rules [Haag, 1992].

Beyond his mathematical axiomatization of quantum mechanics, which (along with its subsequent extension by the C^* -algebraic formulation) lies at the basis of all serious efforts to relate classical and quantum mechanics, von Neumann contributed to this relationship through his analysis of the measurement problem.⁴⁷ Since here the apparent clash between classical and quantum physics comes to a head, it is worth summarizing von Neumann's analysis of this problem here. See also [Wheeler and Zurek, 1983; Busch *et al.*, 1991; Auletta, 2001] and [Mittelstaedt, 2004] for general discussions of the measurement problem.

The essence of the measurement problem is that certain states are never seen in nature, although they are not merely allowed by quantum mechanics (on the assumption of its universal validity), but are even predicted to arise in typical measurement situations. Consider a system S , whose pure states are mathematically described by normalized vectors (more precisely, rays) in a Hilbert space \mathcal{H}_S . One wants to measure an observable \mathcal{O} , which is mathematically represented by a self-adjoint operator O on \mathcal{H}_S . Von Neumann assumes that O has discrete

⁴⁴I.e., when R is self-adjoint with positive spectrum, or, equivalently, when $R = S^*S$ for some $S \in \mathcal{B}(\mathcal{H})$.

⁴⁵Where the 1 in $\text{Exp}(1)$ is the unit operator on \mathcal{H} .

⁴⁶This result has been widely misinterpreted (apparently also by von Neumann himself) as a theorem excluding hidden variables in quantum mechanics. See [Scheibe, 1991]. However, Bell's characterization of von Neumann's linearity assumption in the definition of a state as "silly" is far off the mark, since it holds both in classical mechanics and in quantum mechanics. Indeed, von Neumann's theorem *does* exclude all hidden variable extensions of quantum mechanics that are classical in nature, and it is precisely such extensions that many physicists were originally looking for. See [Rédei and Stöltzner, 2001] and [Scheibe, 2001] for recent discussions of this issue.

⁴⁷Von Neumann [1932] refrained from discussing either the classical limit of quantum mechanics or (probably) the notion of quantization. In the latter direction, he declares that 'If the quantity \mathfrak{R} has the operator R , then the quantity $f(\mathfrak{R})$ has the operator $f(R)$ ', and that 'If the quantities $\mathfrak{R}, \mathfrak{S}, \dots$ have the operators R, S, \dots , then the quantity $\mathfrak{R} + \mathfrak{S} + \dots$ has the operator $R + S + \dots$ '. However, despite his legendary clarity and precision, von Neumann is rather vague about the meaning of the transition $\mathfrak{R} \mapsto R$. It is tempting to construe \mathfrak{R} as a classical observable whose quantum-mechanical counterpart is R , so that the above quotations might be taken as axioms for quantization. However, such an interpretation is neither supported by the surrounding text, nor by our current understanding of quantization (cf. Section 4). For example, a quantization map $\mathfrak{R} \mapsto \mathcal{Q}_h(\mathfrak{R})$ cannot satisfy $f(\mathfrak{R}) \mapsto f(\mathcal{Q}_h(\mathfrak{R}))$ even for very reasonable functions such as $f(x) = x^2$.

spectrum, a simplification which does not hide the basic issues in the measurement problem. Hence O has unit eigenvectors Ψ_n with real eigenvalues o_n . To measure O , one couples the system to an apparatus A with Hilbert space \mathcal{H}_A and “pointer” observable \mathcal{P} , represented by a self-adjoint operator P on \mathcal{H}_A , with discrete eigenvalues p_n and unit eigenvectors Φ_n . The pure states of the total system $S + A$ then correspond to unit vectors in the tensor product $\mathcal{H}_S \otimes \mathcal{H}_A$. A good (“first kind”) measurement is then such that after the measurement, Ψ_n is correlated to Φ_n , that is, for a suitably chosen initial state $I \in \mathcal{H}_A$, a state $\Psi_n \otimes I$ (at $t = 0$) almost immediately evolves into $\Psi_n \otimes \Phi_n$. This can indeed be achieved by a suitable Hamiltonian.

The problem, highlighted by Schrödinger’s cat, now arises if one selects the initial state of S to be $\sum_n c_n \Psi_n$ (with $\sum |c_n|^2 = 1$), for then the superposition principle leads to the conclusion that the final state of the coupled system is $\sum_n c_n \Psi_n \otimes \Phi_n$. Now, basically all von Neumann said was that if one restricts the final state to the system S , then the resulting density matrix is the mixture $\sum_n |c_n|^2 [\Psi_n]$ (where $[\Psi]$ is the orthogonal projection onto a unit vector Ψ),⁴⁸ so that, *from the perspective of the system alone*, the measurement appears to have caused a transition from the pure state $\sum_{n,m} c_n \bar{c}_m \Psi_n \Psi_m^*$ to the mixed state $\sum_n |c_n|^2 [\Psi_n]$, in which interference terms $\Psi_n \Psi_m^*$ for $n \neq m$ are absent. Here the operator $\Psi_n \Psi_m^*$ is defined by $\Psi_n \Psi_m^* f = (\Psi_m, f) \Psi_n$; in particular, $\Psi \Psi^* = [\Psi]$.⁴⁹ Similarly, the apparatus, taken by itself, has evolved from the pure state $\sum_{n,m} c_n \bar{c}_m \Phi_n \Phi_m^*$ to the mixed state $\sum_n |c_n|^2 [\Phi_n]$. This is simply a mathematical theorem (granted the possibility of coupling the system to the apparatus in the desired way), rather than a proposal that there exist two different time-evolutions in Nature, viz. the unitary propagation according to the Schrödinger equation side by side with the above “collapse” process.

In any case, by itself this move by no means solves the measurement problem.⁵⁰ Firstly, in the given circumstances one is not allowed to adopt the ignorance interpretation of mixed states (i.e. assume that the system really is in one of the states Ψ_n); cf., e.g., [Mittelstaedt, 2004]. Secondly, even if one were allowed to do so, one could restore the problem (i.e. the original superposition $\sum_n c_n \Psi_n \otimes \Phi_n$) by once again taking the other component of the system into account.

Von Neumann was well aware of at least this second point, to which he responded by his construction of a *chain*: one redefines $S + A$ as the system, and couples it to a new apparatus B , etc. This eventually leads to a post-measurement state $\sum_n c_n \Psi_n \otimes \Phi_n \otimes \chi_n$ (in hopefully self-explanatory notation, assuming the vectors χ_n form an orthonormal set), whose restriction to $S + A$ is the mixed state $\sum_n |c_n|^2 [\Psi_n] \otimes [\Phi_n]$. The restriction of the latter state to S is, once again, $\sum_n |c_n|^2 [\Psi_n]$. This procedure may evidently be iterated; the point of the construction is evidently to pass on superpositions in some given system to arbitrary

⁴⁸I.e., $[\Psi]f = (\Psi, f)\Psi$; in Dirac notation one would have $[\Psi] = |\Psi\rangle\langle\Psi|$.

⁴⁹In Dirac notation one would have $\Psi_n \Psi_m^* = |\Psi_n\rangle\langle\Psi_m|$.

⁵⁰Not even in an ensemble-interpretation of quantum mechanics, which was the interpretation von Neumann unfortunately adhered to when he wrote his book.

systems higher up in the chain. It follows that for the final state of the original system it does not matter where one “cuts the chain” (that is, which part of the chain one leaves out of consideration), as long as it is done *somewhere*. Von Neumann ([1932], in beautiful prose) and others suggested identifying the cutting with the act of observation, but it is preferable and much more general to simply say that *some* end of the chain is omitted in the description.

The burden of the measurement problem, then, is to

1. Construct a suitable chain along with an appropriate cut thereof; it doesn't matter where the cut is made, as long as it is done.
2. Construct a suitable time-evolution accomplishing the measurement.
3. Justify the ignorance interpretation of mixed states.

As we shall see, these problems are addressed, in a conceptually different but mathematically analogous way, in the Copenhagen interpretation as well as in the decoherence approach. (The main conceptual difference will be that the latter aims to solve also the more ambitious problem of explaining the appearance of the classical world, which in the former seems to be taken for granted).

We conclude this section by saying that despite some brilliant ideas, the founders of quantum mechanics left wide open the problem of deriving classical mechanics as a certain regime of their theory.

3 COPENHAGEN: A REAPPRAISAL

The so-called “Copenhagen interpretation” of quantum mechanics goes back to ideas first discussed and formulated by Bohr, Heisenberg, and Pauli around 1927. Against the idea that there has been a “party line” from the very beginning, it has frequently been pointed out that in the late 1920s there were actually sharp differences of opinion between Bohr and Heisenberg on the interpretation of quantum mechanics and that they never really arrived at a joint doctrine [Hooker, 1972; Stapp, 1972; Hendry, 1984; Beller, 1999; Howard, 2004; Camilleri, 2005]. For example, they never came to agree about the notion of complementarity (see Subsection 3.3). More generally, Heisenberg usually based his ideas on the mathematical formalism of quantum theory, whereas Bohr's position was primarily philosophically oriented. Nonetheless, there is a clearly identifiable core of ideas on which they *did* agree, and since this core has everything to do with the relationship between classical and quantum, we are going to discuss it in some detail.

The principal primary sources are Bohr's Como Lecture, his reply to EPR, and his essay dedicated to Einstein [Bohr, 1927; 1935; 1949].⁵¹ Historical discussions of the emergence and reception of these papers are given in Bohr [1985; 1996] and in Mehra and Rechenberg [2001]. As a selection of the enormous literature

⁵¹These papers were actually written in collaboration with Pauli (after first attempts with Klein), Rosenfeld, and Pais, respectively.

these papers have given rise to, we mention among relatively recent works [Hooker, 1972; Scheibe, 1973; Folse, 1985; Murdoch, 1987; Lahti and Mittelstaedt, 1987; Honner, 1987; Chevalley, 1991; 1999; Faye, 1991; Faye and Folse, 1994; Held, 1994; Howard, 1994; Beller, 1999; Faye, 2002], and [Saunders, 2005]. For Bohr's sparring partners see [Heisenberg, 1930; 1942; 1958; 1984a,b; 1985] with associated secondary literature [Heelan, 1965; Hörz, 1968; Geyer *et al.*, 1993; Camilleri, 2005] and [Pauli, 1933; 1949; 1979; 1985; 1994]), along with [Laurikainen, 1988] and [Enz, 2002].

As with Wittgenstein (and many other thinkers), it helps to understand Bohr if one makes a distinction between an “early” Bohr and a “later” Bohr.⁵² Despite a good deal of continuity in his thought (see below), the demarcation point is his response to EPR [Bohr, 1935],⁵³ and the main shift he made afterwards lies in his sharp insistence on the indivisible unity of object and observer after 1935, focusing on the concept of a *phenomenon*. Before EPR, Bohr equally well believed that object and observer were both necessary ingredients of a complete description of quantum theory, but he then thought that although their interaction could never be neglected, they might at least logically be considered separately. After 1935, Bohr gradually began to claim that object and observer no longer even had separate identities, together forming a “phenomenon”. Accordingly, also his notion of complementarity changed, increasingly focusing on the idea that the specification of the experimental conditions is crucial for the unambiguous use of (necessarily) classical concepts in quantum theory [Scheibe, 1973; Held, 1994]. See also Subsection 3.3 below. This development culminated in Bohr's eventual denial of the existence of the quantum world:

‘There is no quantum world. There is only an abstract quantum-physical description. It is wrong to think that the task of physics is to find out how nature is. Physics concerns what we can say about nature. (...) What is it that we humans depend on? We depend on our words. Our task is to communicate experience and ideas to others. We are suspended in language.’ (quoted by [Petersen, 1963, 8].)⁵⁴

3.1 *The doctrine of classical concepts*

Despite this shift, it seems that Bohr stuck to one key thought throughout his career:

⁵²Here we side with Held [1994] and Beller [1999] against Howard [1994] and Saunders [2005]. See also [Pais, 2000, 22]: ‘Bohr's Como Lecture did not bring the house down, however. He himself would later frown on expressions he used there, such as “disturbing the phenomena by observation”. Such language may have contributed to the considerable confusion that for so long has reigned around this subject.’

⁵³This response is problematic, as is EPR itself. Consequently, there exists a considerable exegetical literature on both, marked by the fact that equally competent and well-informed pairs of commentators manage to flatly contradict each other while at the same time both claiming to explain or reconstruct what Bohr “really” meant.

⁵⁴See [Mermin, 2004] for a witty discussion of this controversial quotation.

‘However far the phenomena transcend the scope of classical physical explanation, the account of all evidence must be expressed in classical terms. (...) The argument is simply that by the word *experiment* we refer to a situation where we can tell others what we have done and what we have learned and that, therefore, the account of the experimental arrangements and of the results of the observations must be expressed in unambiguous language with suitable application of the terminology of classical physics.’ [Bohr, 1949, 209].

This is, in a nutshell, Bohr’s *doctrine of classical concepts*. Although his many drawings and stories may suggest otherwise, Bohr does not quite express the idea here that the goal of physics lies in the description of experiments.⁵⁵ In fact, he merely points out the need for “unambiguous” communication, which he evidently felt threatened by quantum mechanics.⁵⁶ The controversial part of the quote lies in his identification of the means of unambiguous communication with the language of classical physics, involving particles and waves and the like. We will study Bohr’s specific argument in favour of this identification shortly, but it has to be said that, like practically all his foundational remarks on quantum mechanics, Bohr presents his reasoning as self-evident, necessary, and not in need of any further analysis [Scheibe, 1973; Beller, 1999]. Nonetheless, young Heisenberg clashed with Bohr on precisely this point, for Heisenberg felt that the abstract mathematical formalism of quantum theory (rather than Bohr’s world of words and pictures) provided those means of unambiguous communication.⁵⁷

By classical physics Bohr undoubtedly meant the theories of Newton, Maxwell, and Lorentz, but that is not the main point.⁵⁸ For Bohr, the *defining* property of classical physics was the property that it was *objective*, i.e. that it could be studied in an observer-independent way:

‘All description of experiences so far has been based on the assumption, already inherent in ordinary conventions of language, that it is possible to distinguish sharply between the behaviour of objects and the means of observation. This assumption is not only fully justified by everyday experience, *but even constitutes the whole basis of classical physics*’

⁵⁵Which often but misleadingly has been contrasted with Einstein’s belief that the goal of physics is rather to describe reality. See [Landsman, 2006b] for a recent discussion.

⁵⁶Here “unambiguous” means “objective” [Scheibe, 1973; Chevalley, 1991].

⁵⁷It is hard to disagree with Beller’s [1999] conclusion that Bohr was simply not capable of understanding the formalism of post-1925 quantum mechanics, turning his own need of understanding this theory in terms of words and pictures into a deep philosophical necessity.

⁵⁸Otherwise, one should wonder why one shouldn’t use the physics of Aristotle and the scholastics for this purpose, which is a much more effective way of communicating our naive impressions of the world. In contrast, the essence of physics since Newton has been to unmask a reality behind the phenomena. Indeed, Newton himself emphasized that his physics was intended for those capable of natural philosophy, in contrast to *ye vulgar* who believed naive appearances. The fact that Aristotle’s physics is now known to be wrong should not suffice to disqualify its use for Bohr’s purposes, since the very same comment may be made about the physics of Newton etc.

[Bohr, 1958, 25; italics added].⁵⁹

See also [Hooker, 1972; Scheibe, 1973] and [Howard, 1994]. Heisenberg [1958, 55] shared this view:⁶⁰

‘In classical physics science started from the belief — or should one say from the illusion? — that we could describe the world or at least part of the world without any reference to ourselves. This is actually possible to a large extent. We know that the city of London exists whether we see it or not. It may be said that classical physics is just that idealization in which we can speak about parts of the world without any reference to ourselves. Its success has led to the general idea of an objective description of the world.’

On the basis of his “quantum postulate” (see Subsection 3.2), Bohr came to believe that, similarly, the *defining* property of quantum physics was precisely the opposite, i.e. the necessity of the role of the observer (or apparatus — Bohr did not distinguish between the two and never assigned a special role to the mind of the observer or endorsed a subjective view of physics). Identifying unambiguous communication with an objective description, in turn claimed to be the essence of classical physics, Bohr concluded that despite itself quantum physics had to be described entirely in terms of classical physics. Thus his doctrine of classical concepts has an epistemological origin, arising from an analysis of the conditions for human knowledge.⁶¹ In that sense it may be said to be Kantian in spirit [Hooker, 1972; Murdoch, 1987; Chevalley, 1991; 1999].

Now, Bohr himself is on record as saying: ‘They do it smartly, but what counts is to do it right’ [Rosenfeld, 1967, 129].⁶² The doctrine of classical concepts is certainly smart, but is it right? As we have seen, Bohr’s argument starts from the claim that classical physics is objective (or ‘unambiguous’) in being independent of the observer. In fact, nowadays it is widely believed that quantum mechanics leads to the *opposite* conclusion that “quantum reality” (whatever that may be) is objective (though “veiled” in the terminology of d’Espagnat [1995]), while “classical reality” only comes into existence relative to a certain specification: this is

⁵⁹Despite the typical imperative tone of this quotation, Bohr often regarded certain other properties as essential to classical physics, such as determinism, the combined use of space-time concepts and dynamical conservation laws, and the possibility of pictorial descriptions. However, these properties were in some sense secondary, as Bohr considered them to be *consequences* of the possibility of isolating an object in classical physics. For example: ‘The assumption underlying the ideal of causality [is] that the behaviour of the object is uniquely determined, quite independently of whether it is observed or not’ [Bohr, 1937], and then again, now negatively: ‘the renunciation of the ideal of causality [in quantum mechanics] is founded logically only on our not being any longer in a position to speak of the autonomous behaviour of a physical object’ [Bohr, 1937]. See [Scheibe, 1973].

⁶⁰As Camilleri [2005, 161] states: ‘For Heisenberg, classical physics is the fullest expression of the ideal of objectivity.’

⁶¹See, for example, the very *title* of Bohr [1958]!

⁶²‘They’ refers to EPR.

stance 1 discussed in the Introduction.⁶³ Those who disagree with stance 1 cannot use stance 2 (of denying the fundamental nature of quantum theory) at this point either, as that is certainly not what Bohr had in mind. Unfortunately, in his most outspoken defence of Bohr, even Heisenberg [1958, 55] was unable to find a better argument for Bohr's doctrine than the lame remark that 'the use of classical concepts is finally a consequence of the general human way of thinking.'⁶⁴

In our opinion, Bohr's *motivation* for his doctrine has to be revised in the light of our current understanding of quantum theory; we will do so in Subsection 3.4. In any case, whatever its motivation, the doctrine *itself* seems worth keeping: apart from the fact that it evidently describes experimental practice, it provides a convincing explanation for the probabilistic nature of quantum mechanics (cf. the next subsection).

3.2 *Object and apparatus: the Heisenberg cut*

Describing quantum physics in terms of classical concepts sounds like an impossible and even self-contradictory task (cf. [Heisenberg, 1958]). For one, it precludes a completely quantum-mechanical description of the world: 'However far the phenomena transcend the scope of classical physical explanation, the account of all evidence must be expressed in classical terms.' But at the same time it precludes a purely classical description of the world, for underneath classical physics one has quantum theory.⁶⁵ The fascination of Bohr's philosophy of quantum mechanics lies precisely in his brilliant resolution of this apparently paradoxical situation.

The first step of this resolution that he and Heisenberg proposed is to divide the system whose description is sought into two parts: one, the object, is to be described quantum-mechanically, whereas the other, the apparatus, is treated *as if it were classical*. Despite innumerable claims to the contrary in the literature (i.e. to the effect that Bohr held that a separate realm of Nature was intrinsically classical), there is no doubt that both Bohr and Heisenberg believed in the fundamental and universal nature of quantum mechanics, and saw the classical description of the apparatus as a *purely epistemological move without any coun-*

⁶³Indeed, interesting recent attempts to make Bohr's philosophy of quantum mechanics precise accommodate the a priori status of classical observables into some version of the modal interpretation; see [Dieks, 1989b; Bub, 1999; Halvorson and Clifton, 1999; 2002], and [Dickson, 2005]. It should give one some confidence in the possibility of world peace that the two most hostile interpretations of quantum mechanics, viz. Copenhagen and Bohm [Cushing, 1994] have now found a common home in the modal interpretation in the sense of the authors just cited! Whether or not one agrees with Bub's [2004] criticism of the modal interpretation, Bohr's insistence on the necessity of classical concepts is not vindicated by any current version of it.

⁶⁴And similarly: 'We are forced to use the language of classical physics, simply because we have no other language in which to express the results.' [Heisenberg, 1971, 130]. This in spite of the fact that the later Heisenberg thought about this matter very deeply; see, e.g., his [1942], as well as [Camilleri, 2005]. Murdoch [1987, 207–210] desperately tries to boost the doctrine of classical concepts into a profound philosophical argument by appealing to Strawson [1959].

⁶⁵This peculiar situation makes it very hard to give a realist account of the Copenhagen interpretation, since quantum reality is denied whereas classical reality is neither fundamental nor real.

terpart in ontology, expressing the fact that a given *quantum* system is *being used* as a measuring device.⁶⁶ For example: ‘The construction and the functioning of all apparatus like diaphragms and shutters, serving to define geometry and timing of the experimental arrangements, or photographic plates used for recording the localization of atomic objects, will depend on properties of materials which are themselves essentially determined by the quantum of action’ [Bohr, 1948, 315], as well as: ‘We are free to make the cut only within a region where the quantum mechanical description of the process concerned is effectively equivalent with the classical description’ [Bohr, 1935, 701].⁶⁷

The separation between object and apparatus called for here is usually called the *Heisenberg cut*, and it plays an absolutely central role in the Copenhagen interpretation of quantum mechanics.⁶⁸ The idea, then, is that *a quantum-mechanical object is studied exclusively through its influence on an apparatus that is described classically*. Although *described* classically, the apparatus is a quantum system, and is supposed to be influenced by its *quantum-mechanical* coupling to the underlying (quantum) object.

The alleged necessity of including both object and apparatus in the description was initially claimed to be a consequence of the so-called “quantum postulate”. This notion played a key role in Bohr’s thinking: his Como Lecture [Bohr, 1927] was even entitled ‘The quantum postulate and the recent development of atomic theory’. There he stated its contents as follows: ‘The essence of quantum theory is the quantum postulate: every atomic process has an essential discreteness — completely foreign to classical theories — characterized by Planck’s quantum of action.’⁶⁹ Even more emphatically, in his reply to EPR [Bohr, 1935, 697]: ‘Indeed the finite interaction between object and measuring agencies conditioned by the

⁶⁶See especially [Scheibe, 1973] on Bohr, and [Heisenberg, 1958]). The point in question has also been made by R. Haag (who knew both Bohr and Heisenberg) in most of his talks on quantum mechanics in the 1990s. In this respect we disagree with Howard [1994], who claims that according to Bohr a classical description of an apparatus amounts to picking a particular (maximally) abelian subalgebra of its quantum-mechanical algebra of ‘beables’, which choice is dictated by the measurement context. But having a commutative algebra falls far short of a classical description, since in typical examples one obtains only half of the canonical classical degrees of freedom in this way. Finding a classical description of a quantum-mechanical system is a much deeper problem, to which we shall return throughout this paper.

⁶⁷This last point suggests that the cut has something to do with the division between a microscopic and a macroscopic realm in Nature, but although this division often facilitates making the cut when it is well defined, this is by no means a matter of principle. Cf. [Howard, 1994]. In particular, all objections to the Copenhagen interpretation to the effect that the interpretation is ill-defined because the micro-macro distinction is blurred are unfounded.

⁶⁸Pauli [1949] went as far as saying that the Heisenberg cut provides the appropriate generalization modern physics offers of the old Kantian opposition between a knowable object and a knowing subject: ‘Auf diese Weise verallgemeinert die moderne Physik die alte Gegenüberstellung von erkennenden Subjekt auf der einen Seite und des erkannten Objektes auf der anderen Seite zu der Idee des Schnittes zwischen Beobachter oder Beobachtungsmittel und dem beobachteten System.’ (‘In this way, modern physics generalizes the old opposition between the knowing subject on the one hand and the known object on the other to the idea of the cut between observer or means of observation and the observed system.’) He then continued calling the cut a necessary condition for human knowledge: see footnote 72.

⁶⁹Instead of ‘discreteness’, Bohr alternatively used the words ‘discontinuity’ or ‘individuality’

very existence of the quantum of action entails — because of the impossibility of controlling the reaction of the object on the measurement instruments if these are to serve their purpose — the necessity of a final renunciation of the classical ideal of causality and a radical revision of our attitude towards the problem of physical reality.’ Also, Heisenberg’s uncertainty relations were originally motivated by the quantum postulate in the above form. According to Bohr and Heisenberg around 1927, this ‘essential discreteness’ causes an ‘uncontrollable disturbance’ of the object by the apparatus during their interaction. Although the “quantum postulate” is not supported by the mature mathematical formalism of quantum mechanics and is basically obsolete, the intuition of Bohr and Heisenberg that a measurement of a quantum-mechanical object causes an ‘uncontrollable disturbance’ of the latter is actually quite right.⁷⁰

In actual fact, the reason for this disturbance does not lie in the “quantum postulate”, but in the phenomenon of entanglement, as further discussed in Subsection 3.4. Namely, from the point of view of von Neumann’s measurement theory (see Subsection 2.5) the Heisenberg cut is just a two-step example of a von Neumann chain, with the special feature that after the quantum-mechanical interaction has taken place, the second link (i.e. the apparatus) is *described* classically. The latter feature not only supports Bohr’s philosophical agenda, but, more importantly, also suffices to guarantee the applicability of the ignorance interpretation of the mixed state that arises after completion of the measurement.⁷¹ All of von Neumann’s analysis of the arbitrariness of the location of the cut applies here, for one may always extend the definition of the quantum-mechanical object by coupling the original choice to any other purely quantum-mechanical system one likes, and analogously for the classical part. Thus the two-step nature of the Heisenberg cut includes the possibility that the first link or object is in fact a lengthy chain in itself (as long as it is quantum-mechanical), and similarly for the second link (as long as it is classical).⁷² This arbitrariness, subject to the limitation expressed by the second [1935] Bohr quote in this subsection, was well recognized by Bohr and Heisenberg, and was found at least by Bohr to be of great philosophical importance.

It is the interaction between object and apparatus that causes the measurement to ‘disturb’ the former, but it is only and precisely the classical description of

as well. He rarely omitted amplifications like ‘essential’.

⁷⁰Despite the fact that Bohr later distanced himself from it; cf. [Beller, 1999] and footnote 52 above. In a correct analysis, what is disturbed upon coupling to a classical apparatus is the quantum-mechanical state of the object (rather than certain sharp values of classical observables such as position and momentum, as the early writings of Bohr and Heisenberg suggest).

⁷¹In a purely quantum-mechanical von Neumann chain the final state of system plus apparatus is pure, but if the apparatus is classical, then the post-measurement state is mixed.

⁷²Pauli [1949] once more: ‘Während die EXISTENZ eines solchen Schnittes eine notwendige Bedingung menschlicher Erkenntnis ist, faßt sie die LAGE des Schnittes als bis zu einem gewissen Grade willkürlich und als Resultat einer durch Zweckmäßigkeitserwägungen mitbestimmten, also teilweise freien Wahl auf.’ (‘While the EXISTENCE of such a [Heisenberg] cut is a necessary condition for human knowledge, its LOCATION is to some extent arbitrary as a result of a pragmatic and thereby partly free choice.’)

the latter that (through the ignorance interpretation of the final state) makes the disturbance ‘uncontrollable’.⁷³ In the Copenhagen interpretation, *probabilities arise solely because we look at the quantum world through classical glasses.*

‘Just the necessity of accounting for the function of the measuring agencies on classical lines excludes in principle in proper quantum phenomena an accurate control of the reaction of the measuring instruments on the atomic objects.’ [Bohr, 1956, 87]

‘One may call these uncertainties objective, in that they are simply a consequence of the fact that we describe the experiment in terms of classical physics; they do not depend in detail on the observer. One may call them subjective, in that they reflect our incomplete knowledge of the world.’ [Heisenberg, 1958, 53–54]

Thus the picture that arises is this: Although the quantum-mechanical side of the Heisenberg cut is described by the Schrödinger equation (which is deterministic), while the classical side is subject to Newton’s laws (which are equally well deterministic),⁷⁴ unpredictability arises because the quantum system serving as an apparatus is approximated by a classical system. The ensuing probabilities reflect the ignorance arising from the decision (or need) to ignore the quantum-mechanical degrees of freedom of the apparatus. Hence the probabilistic nature of quantum theory is not intrinsic but extrinsic, and as such is entirely a consequence of the doctrine of classical concepts, which by the same token *explains* this nature.

Mathematically, the simplest illustration of this idea is as follows. Take a finite-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^n$ with the ensuing algebra of observables $\mathcal{A} = M_n(\mathbb{C})$ (i.e. the $n \times n$ matrices). A unit vector $\Psi \in \mathbb{C}^n$ determines a quantum-mechanical state in the usual way. Now describe this quantum system as if it were classical by ignoring all observables except the diagonal matrices. The state then immediately collapses to a probability measure on the set of n points, with probabilities given by the Born rule $p(i) = |(e_i, \Psi)|^2$, where $(e_i)_{i=1, \dots, n}$ is the standard basis of \mathbb{C}^n . Similarly, the Born–Pauli rule for the probabilistic interpretation of the wave function $\Psi \in L^2(\mathbb{R}^3)$ in terms of $|\Psi(x)|^2$ immediately follows if one ignores all observables on $L^2(\mathbb{R}^3)$ except the position operator.⁷⁵

Despite the appeal of this entire picture, it is not at all clear that it actually applies! There is no a priori guarantee whatsoever that one may indeed describe a quantum system “as if it were classical”. Bohr and Heisenberg apparently took the existence of the classical world of mountains and creeks they loved so much for granted, the former probably on empirical grounds, the latter on the basis of his

⁷³These points were not clearly separated by Heisenberg [1927] in his paper on the uncertainty relations, but were later clarified by Bohr. See [Scheibe, 1973].

⁷⁴But see [Earman, 1986; 2005].

⁷⁵Technically, one restricts Ψ — seen as a state on the C^* -algebra $\mathcal{B}(L^2(\mathbb{R}^3))$ — to the C^* -algebra $C_0(\mathbb{R}^3)$ given by all multiplication operators on $L^2(\mathbb{R}^3)$ defined by continuous functions of $x \in \mathbb{R}^3$ that vanish at infinity. This restriction yields a probability measure on \mathbb{R}^3 , which is precisely the usual one originally proposed by Pauli.

own uncertainty relations — both almost blind to the extremely delicate mathematical and conceptual problems involved. In our opinion, the main difficulty in making sense of the Copenhagen interpretation therefore lies in the justification of the classical description in question. This difficulty is the main topic of this paper, of which Section 6 is of particular relevance in the present context.

3.3 Complementarity

The notion of a Heisenberg cut is subject to a certain arbitrariness even apart from the precise location of the cut within a given chain, for one might in principle construct the chain in various different and incompatible ways. This arbitrariness was analyzed by Bohr in terms of what he called *complementarity*.⁷⁶

Bohr never gave a precise definition of complementarity,⁷⁷ but restricted himself to the analysis of a number of examples.⁷⁸ A prominent such example is the complementarity between a “causal”⁷⁹ description of a quantum system in which conservation laws hold, and a space-time description that is necessarily statistical in character. Here Bohr’s idea seems to have been that a stationary state (i.e. an energy eigenstate) of an atom is incompatible with an electron moving in its orbit in space and time — see Subsection 5.4 for a discussion of this issue. Heisenberg [1958], however, took this example of complementarity to mean that a system on which no measurement is performed evolves deterministically according to the Schrödinger equation, whereas a rapid succession of measurements produces a space-time path whose precise form quantum theory is only able to predict statistically [Camilleri, 2005]. In other words, this example reproduces precisely the picture through which Heisenberg [1927] believed he had established the connection between classical and quantum mechanics; cf. Subsection 2.3.

Bohr’s other key example was the complementarity between particles and waves. Here his principal aim was to make sense of Young’s double-slit experiment. The well-known difficulty with a classical visualization of this experiment is that a particle description appears impossible because a particle has to go through a single slit, ruining the interference pattern gradually built up on the detection screen, whereas a wave description seems incompatible with the point-like localization on the screen once the wave hits it. Thus Bohr suggested that whilst each of these classical descriptions is incomplete, the union of them is necessary for a complete

⁷⁶Unfortunately and typically, Bohr once again presented complementarity as a necessity of thought rather than as the truly amazing possible mode of description it really is.

⁷⁷Perhaps he preferred this approach because he felt a definition could only reveal part of what was supposed to be defined: one of his favourite examples of complementarity was that between definition and observation.

⁷⁸We refrain from discussing the complementarity between truth and clarity, science and religion, thoughts and feelings, and objectivity and introspection here, despite the fact that on this basis Bohr’s biographer Pais [1997] came to regard his subject as the greatest philosopher since Kant.

⁷⁹Bohr’s use the word “causal” is quite confusing in view of the fact that in the British empiricist tradition causality is often interpreted in the sense of a space-time description. But Bohr’s “causal” is meant to be *complementary* to a space-time description!

description of the experiment.

The deeper epistemological point appears to be that although the *completeness* of the quantum-mechanical description of the microworld systems seems to be endangered by the doctrine of classical concepts, it is actually restored by the inclusion of *two* “complementary” descriptions (i.e. of a given quantum system plus a measuring device that is necessarily described classically, ‘if it is to serve its purpose’). Unfortunately, despite this attractive general idea it is unclear to what precise definition of complementarity Bohr’s examples should lead. In the first, the complementary notions of determinism and a space-time description are in mutual harmony as far as classical physics is concerned, but are apparently in conflict with each other in quantum mechanics. In the second, however, the wave description of some entity contradicts a particle description of the same entity precisely in classical physics, whereas in quantum mechanics these descriptions somehow coexist.⁸⁰

Scheibe [1973, 32] notes a ‘clear convergence [in the writings of Bohr] towards a preferred expression of a complementarity between phenomena’, where a Bohrian *phenomenon* is an indivisible union (or “whole”) of a quantum system and a classically described experimental arrangement used to study it; see item 2 below. Some of Bohr’s early examples of complementarity can be brought under this heading, others cannot [Held, 1994]. For many students of Bohr (including the present author), the fog has yet to clear up.⁸¹ Nonetheless, the following mathematical interpretations might assign some meaning to the idea of complementarity in the framework of von Neumann’s formalism of quantum mechanics.⁸²

1. Heisenberg [1958] *identified complementary pictures of a quantum-mechanical system with equivalent mathematical representations thereof*. For example, he thought of the complementarity of x and p as the existence of what we now call the Schrödinger representations of the canonical commutation relations (CCR) on $L^2(\mathbb{R}^n)$ and its Fourier transform to momentum space. Furthermore, he felt that in quantum field theory particles and waves gave two *equivalent* modes of description of quantum theory because of second quantization. Thus for Heisenberg complementary pictures are classical because there is an underlying classical variable, with no apparatus in sight, and such pictures are not mutually contradictory but (unitarily) equivalent.

⁸⁰On top of this, Bohr mixed these examples in conflicting ways. In discussing bound states of electrons in an atom he jointly made determinism and particles one half of a complementary pair, waves and space-time being the other. In his description of electron-photon scattering he did it the other way round: this time determinism and waves formed one side, particles and space-time the other (cf. [Beller, 1999]).

⁸¹Even Einstein [1949, 674] conceded that throughout his debate with Bohr he had never understood the notion of complementarity, ‘the sharp formulation of which, moreover, I have been unable to achieve despite much effort which I have expended on it.’ See [Landsman, 2006b] for the author’s view on the Bohr–Einstein debate.

⁸²This exercise is quite against the spirit of Bohr, who is on record as saying that ‘von Neumann’s approach (...) did not *solve* problems but created *imaginary difficulties* ([Scheibe, 1973, 11], quoting Feyerabend; italics in original).

See also [Camilleri, 2005, 88], according to whom ‘Heisenberg never accepted Bohr’s complementarity arguments’.

2. Pauli [1933] simply stated that *two observables are complementary when the corresponding operators fail to commute*.⁸³ Consequently, it then follows from Heisenberg’s uncertainty relations that complementary observables cannot be measured simultaneously with arbitrary precision. This suggests (but by no means proves) that they should be measured independently, using mutually exclusive experimental arrangements. The latter feature of complementarity was emphasized by Bohr in his later writings.⁸⁴ This approach makes the notion of complementarity unambiguous and mathematically precise, and perhaps for this reason the few physicists who actually use the idea of complementarity in their work tend to follow Pauli and the later Bohr.⁸⁵
3. The present author proposes that *observables and pure states are complementary*. For in the Schrödinger representation of elementary quantum mechanics, the former are, roughly speaking, generated by the position and

⁸³More precisely, one should probably require that the two operators in question generate the ambient algebra of observables, so that complementarity in Pauli’s sense is really defined between two commutative subalgebras of a given algebra of observables (again, provided they jointly generate the latter).

⁸⁴Bohr’s earlier writings do not quite conform to Pauli’s approach [Scheibe, 1973; Held, 1994]. In Bohr’s discussions of the double-slit experiment particle and wave form a complementary pair, whereas Pauli’s complementary observables are position and momentum, which refer to a single side of Bohr’s pair. For a precise analysis of the relationship between mutually exclusive experimental arrangements, non-commuting observables, and the Heisenberg uncertainty relations see [Busch *et al.*, 1998] and [De Muynck, 2002].

⁸⁵Adopting this point of view, it is tempting to capture the complementarity between position and momentum by means of the following conjecture: *Any normal pure state ω on $\mathcal{B}(L^2(\mathbb{R}^n))$ (that is, any wave function seen as a state in the sense of C^* -algebras) is determined by the pair $\{\omega|L^\infty(\mathbb{R}^n), \omega|FL^\infty(\mathbb{R}^n)F^{-1}\}$ (in other words, by its restrictions to position and momentum). Here $L^\infty(\mathbb{R}^n)$ is the von Neumann algebra of multiplication operators on $L^2(\mathbb{R}^n)$, i.e. the von Neumann algebra generated by the position operator, whereas $FL^\infty(\mathbb{R}^n)F^{-1}$ is its Fourier transform, i.e. the von Neumann algebra generated by the momentum operator. The idea is that each of its restrictions $\omega|L^\infty(\mathbb{R}^n)$ and $\omega|FL^\infty(\mathbb{R}^n)F^{-1}$ gives a classical picture of ω . These restrictions are a measure on \mathbb{R}^n interpreted as position space, and another measure on \mathbb{R}^n interpreted as momentum space. Unfortunately, this conjecture is false. The following counterexample was provided by D. Buchholz (private communication): take ω as the state defined by the wave function $\Psi(x) \sim \exp(-ax^2/2)$ with $\text{Re}(a) > 0$, $\text{Im}(a) \neq 0$, and $|a|^2 = 1$. Then ω depends on $\text{Im}(a)$, whereas neither $\omega|L^\infty(\mathbb{R}^n)$ nor $\omega|FL^\infty(\mathbb{R}^n)F^{-1}$ does. There is even a counterexample to the analogous conjecture for the C^* -algebra of 2×2 matrices, found by H. Halvorson: if A is the commutative C^* -algebra generated by σ_x , and B the one generated by σ_y , then the two different eigenstates of σ_z coincide on A and on B . One way to improve our conjecture might be to hope that if, in the Schrödinger picture, two states coincide on the two given commutative von Neumann algebras for all times, then they must be equal. But this can only be true for certain “realistic” time-evolutions, for the trivial Hamiltonian $H = 0$ yields the above counterexample. We leave this as a problem for future research. At the time of writing, Halvorson [2004] contains the only sound mathematical interpretation of the complementarity between position and momentum, by relating it to the representation theory of the CCR. He shows that in any representation where the position operator has eigenstates, there is no momentum operator, and vice versa.*

momentum operators, whereas the latter are given by wave functions. Some of Bohr's other examples of complementarity also square with this interpretation (at least if one overlooks the collapse of the wavefunction upon a measurement). Here one captures the idea that both ingredients of a complementary pair are necessary for a complete description, though the alleged mutual contradiction between observables and states is vague. Also, this reading of complementarity relies on a specific representation of the canonical commutation relations. It is not quite clear what one gains with this ideology, but perhaps it deserves to be developed in some more detail. For example, in quantum field theory it is once more the observables that carry the space-time description, especially in the algebraic description of Haag [1992].

3.4 *Epilogue: entanglement to the rescue?*

Bohr's "quantum postulate" being obscure and obsolete, it is interesting to consider Howard's [1994] 'reconstruction' of Bohr's philosophy of physics on the basis of entanglement.⁸⁶ His case can perhaps be strengthened by an appeal to the analysis Primas [1983] has given of the need for classical concepts in quantum physics.⁸⁷ Primas proposes to define a "quantum object" as a physical system \mathcal{S} that is free from what he calls "EPR-correlations" with its environment. Here the "environment" is meant to include apparatus, observer, the rest of the universe if necessary, and what not. In elementary quantum mechanics, quantum objects in this sense exist only in very special states: if \mathcal{H}_S is the Hilbert space of the system S , and \mathcal{H}_E that of the environment E , any pure state of the form $\sum_i c_i \Psi_i \otimes \Phi_i$ (with more than one term) by definition correlates S with E ; the only uncorrelated pure states are those of the form $\Psi \otimes \Phi$ for unit vectors $\Psi \in \mathcal{H}_S$, $\Phi \in \mathcal{H}_E$. The restriction of a pure EPR-correlated state on $S + E$ to S is mixed, so that the (would-be) quantum object 'does not have its own pure state'; in other words, the restriction of an EPR-correlated state ω to S together with its restriction to E do not jointly determine ω . More generally, if the state of $S + E$ is EPR-correlated, a complete characterization of the state of S requires E (and vice versa). But (against Bohr!) mathematics defeats words: the sharpest characterization of the notion of EPR-correlations can be given in terms of operator algebras, as follows. In the spirit of the remainder of the paper we proceed in a rather general and abstract way.⁸⁸ For what follows see especially [Werner, 1989].

⁸⁶We find little evidence that Bohr himself ever thought along those lines. With approval we quote Zeh, who, following a statement of the quantum postulate by Bohr similar to the one in Subsection 3.2 above, writes: 'The later revision of these early interpretations of quantum theory (required by the important role of entangled quantum states for much larger systems) seems to have gone unnoticed by many physicists.' [Joos *et al.*, 2003, 23] See also [Howard, 1990] for an interesting historical perspective on entanglement, and cf. [Raimond *et al.*, 2001] for the experimental situation.

⁸⁷See also [Amann and Primas, 1997] and [Primas, 1997].

⁸⁸Summers and Werner [1987] give even more general results, where the tensor product $\mathcal{A} \hat{\otimes} \mathcal{B}$ below is replaced by an arbitrary C^* -algebra \mathcal{C} containing \mathcal{A} and \mathcal{B} as C^* -subalgebras.

Let \mathcal{A} and \mathcal{B} be C^* -algebras,⁸⁹ with tensor product $\mathcal{A} \hat{\otimes} \mathcal{B}$.⁹⁰ Less abstractly, just think of two Hilbert spaces \mathcal{H}_S and \mathcal{H}_E as above, with tensor product $\mathcal{H}_S \otimes \mathcal{H}_E$, and assume that $\mathcal{A} = \mathcal{B}(\mathcal{H}_S)$ while \mathcal{B} is either $\mathcal{B}(\mathcal{H}_E)$ itself or some (norm-closed and involutive) commutative subalgebra thereof. The tensor product $\mathcal{A} \hat{\otimes} \mathcal{B}$ is then a (norm-closed and involutive) subalgebra of $\mathcal{B}(\mathcal{H}_S \otimes \mathcal{H}_E)$, the algebra of all bounded operators on $\mathcal{H}_S \otimes \mathcal{H}_E$.

A *product state* on $\mathcal{A} \hat{\otimes} \mathcal{B}$ is a state of the form $\omega = \rho \otimes \sigma$, where the states ρ on \mathcal{A} and σ on \mathcal{B} may be either pure or mixed.⁹¹ We say that a state ω on $\mathcal{A} \hat{\otimes} \mathcal{B}$ is *decomposable* when it is a mixture of product states, i.e. when $\omega = \sum_i p_i \rho_i \otimes \sigma_i$, where the coefficients $p_i > 0$ satisfy $\sum_i p_i = 1$.⁹² A decomposable state ω is pure precisely when it is a product of pure states. This has the important consequence that both its restrictions $\omega|_{\mathcal{A}}$ and $\omega|_{\mathcal{B}}$ to \mathcal{A} and \mathcal{B} , respectively, are pure as well.⁹³ On the other hand, a state on $\mathcal{A} \hat{\otimes} \mathcal{B}$ may be said to be *EPR-correlated* [Primas, 1983] when it is *not* decomposable. An EPR-correlated *pure* state has the property that its restriction to \mathcal{A} or \mathcal{B} is *mixed*.

Raggio [1981] proved that the following two conditions are equivalent:

⁸⁹ Recall that a C^* -algebra is a complex algebra \mathcal{A} that is complete in a norm $\|\cdot\|$ that satisfies $\|AB\| \leq \|A\| \|B\|$ for all $A, B \in \mathcal{A}$, and has an involution $A \rightarrow A^*$ such that $\|A^*A\| = \|A\|^2$. A basic examples is $\mathcal{A} = \mathcal{B}(\mathcal{H})$, the algebra of all bounded operators on a Hilbert space \mathcal{H} , equipped with the usual operator norm and adjoint. By the Gelfand–Naimark theorem, any C^* -algebra is isomorphic to a norm-closed self-adjoint subalgebra of $\mathcal{B}(\mathcal{H})$, for some Hilbert space \mathcal{H} . Another key example is $\mathcal{A} = C_0(X)$, the space of all continuous complex-valued functions on a (locally compact Hausdorff) space X that vanish at infinity (in the sense that for every $\varepsilon > 0$ there is a *compact* subset $K \subset X$ such that $|f(x)| < \varepsilon$ for all $x \notin K$), equipped with the supremum norm $\|f\|_\infty := \sup_{x \in X} |f(x)|$, and involution given by (pointwise) complex conjugation. By the Gelfand–Naimark lemma, any commutative C^* -algebra is isomorphic to $C_0(X)$ for some locally compact Hausdorff space X .

⁹⁰ The tensor product of two (or more) C^* -algebras is not unique, and we here need the so-called *projective* tensor product $\mathcal{A} \hat{\otimes} \mathcal{B}$, defined as the completion of the algebraic tensor product $\mathcal{A} \otimes \mathcal{B}$ in the *maximal* C^* -cross-norm. The choice of the projective tensor product guarantees that each state on $\mathcal{A} \otimes \mathcal{B}$ extends to a state on $\mathcal{A} \hat{\otimes} \mathcal{B}$ by continuity; conversely, since $\mathcal{A} \otimes \mathcal{B}$ is dense in $\mathcal{A} \hat{\otimes} \mathcal{B}$, each state on the latter is uniquely determined by its values on the former. See [Wegge-Olsen, 1993, Appendix T], or [Takesaki, 2003, Vol. I, Ch. IV]. In particular, product states $\rho \otimes \sigma$ and mixtures $\omega = \sum_i p_i \rho_i \otimes \sigma_i$ thereof as considered below are well defined on $\mathcal{A} \hat{\otimes} \mathcal{B}$. If $\mathcal{A} \subset \mathcal{B}(\mathcal{H}_S)$ and $\mathcal{B} \subset \mathcal{B}(\mathcal{H}_E)$ are von Neumann algebras, as in the analysis of Raggio [1981; 1988], it is easier (and sufficient) to work with the *spatial* tensor product $\mathcal{A} \overline{\otimes} \mathcal{B}$, defined as the double commutant (or weak completion) of $\mathcal{A} \otimes \mathcal{B}$ in $\mathcal{B}(\mathcal{H}_S \otimes \mathcal{H}_E)$. For any *normal* state on $\mathcal{A} \otimes \mathcal{B}$ extends to a normal state on $\mathcal{A} \overline{\otimes} \mathcal{B}$ by continuity.

⁹¹We use the notion of a state that is usual in the algebraic framework. Hence a *state* on a C^* -algebra \mathcal{A} is a linear functional $\rho : \mathcal{A} \rightarrow \mathbb{C}$ that is *positive* in that $\rho(A^*A) \geq 0$ for all $A \in \mathcal{A}$ and *normalized* in that $\rho(1) = 1$, where 1 is the unit element of \mathcal{A} . If \mathcal{A} is a von Neumann algebra, one has the notion of a *normal* state, which satisfies an additional continuity condition. If $\mathcal{A} = \mathcal{B}(\mathcal{H})$, then a fundamental theorem of von Neumann states that each normal state ρ on \mathcal{A} is given by a density matrix $\hat{\rho}$ on \mathcal{H} , so that $\rho(A) = \text{Tr}(\hat{\rho}A)$ for each $A \in \mathcal{A}$. In particular, a normal pure state on $\mathcal{B}(\mathcal{H})$ (seen as a von Neumann algebra) is necessarily of the form $\psi(A) = (\Psi, A\Psi)$ for some unit vector $\Psi \in \mathcal{H}$.

⁹²Infinite sums are allowed here. More precisely, ω is decomposable if it is in the w^* -closure of the convex hull of the product states on $\mathcal{A} \hat{\otimes} \mathcal{B}$.

⁹³The restriction $\omega|_{\mathcal{A}}$ of a state ω on $\mathcal{A} \hat{\otimes} \mathcal{B}$ to, say, \mathcal{A} is given by $\omega|_{\mathcal{A}}(A) = \omega(A \otimes 1)$, where 1 is the unit element of \mathcal{B} , etc.

- Each state on $\mathcal{A} \hat{\otimes} \mathcal{B}$ is decomposable;
- \mathcal{A} or \mathcal{B} is commutative.

In other words, EPR-correlated states exist precisely when \mathcal{A} and \mathcal{B} are both non-commutative.⁹⁴ As one might expect, this result is closely related to the Bell inequalities. Namely, consider the inequality

$$(1) \quad |\sup\{\omega(A_1(B_1 + B_2) + A_2(B_1 - B_2))\}| \leq 2,$$

where ω is a fixed state on $\mathcal{A} \hat{\otimes} \mathcal{B}$ and the supremum is taken over all self-adjoint operators $A_1, A_2 \in \mathcal{A}$, $B_1, B_2 \in \mathcal{B}$, each of norm ≤ 1 . Eq. (1) holds if and only if ω is decomposable [Baez, 1987; Raggio, 1988]. Consequently, the inequality (1) can only be violated in some state ω when the algebras \mathcal{A} and \mathcal{B} are both noncommutative. If, on the other hand, (1) is satisfied, then one knows that there exists a classical probability space and probability measure (and hence a “hidden variables” theory) reproducing the given correlations [Pitowsky, 1989]. As stressed by Bacciagaluppi [1993], such a description does *not* require the entire setting to be classical; as we have seen, only one of the algebras \mathcal{A} and \mathcal{B} has to be commutative for the Bell inequalities to hold.

Where does this leave us with respect to Bohr? If we follow Primas [1983] in describing a (quantum) object as a system free from EPR-correlations with its environment, then the mathematical results just reviewed leave us with two possibilities. Firstly, we may pay lip-service to Bohr in taking the algebra \mathcal{B} (interpreted as the algebra of observables of the environment in the widest possible sense, as above) to be commutative *as a matter of description*. In that case, our object is really an “object” in *any* of its states. But clearly this is not the only possibility. For even in the case of elementary quantum mechanics — where $\mathcal{A} = \mathcal{B}(\mathcal{H}_S)$ and $\mathcal{B} = \mathcal{B}(\mathcal{H}_E)$ — the system is still an “object” in the sense of Primas as long as the total state ω of $S + E$ is decomposable. In general, for pure states this just means that $\omega = \psi \otimes \phi$, i.e. that the total state is a product of pure states. To accomplish this, one has to define the Heisenberg cut in an appropriate way, and subsequently hope that the given product state remains so under time-evolution (see [Amann and Primas, 1997] and [Atmanspacher *et al.*, 1999], and references therein). This selects certain states on \mathcal{A} as “robust” or “stable”, in much the same way as in the decoherence approach. We therefore continue this discussion in Section 7 (see especially point 6 in Subsection 7.1).

4 QUANTIZATION

Heisenberg’s [1925] idea of *Umdeutung* (*reinterpretation*) suggests that it is possible to construct a quantum-mechanical description of a physical system whose classical description is known. As we have seen, this possibility was realized by Schrödinger

⁹⁴Raggio [1981] proved this for von Neumann algebras and normal states. His proof was adapted to C^* -algebras by Bacciagaluppi [1993].

[1925c], who found the simplest example (2) and (3) of *Umdeutung* in the context of atomic physics. This early example was phenomenally successful, as almost all of atomic and molecular physics is still based on it.

Quantization theory is an attempt to understand this example, make it mathematically precise, and generalize it to more complicated systems. It has to be stated from the outset that, like the entire classical-quantum interface, the nature of quantization is not yet well understood. This fact is reflected by the existence of a fair number of competing quantization procedures, the most transparent of which we will review below.⁹⁵ Among the first mathematically serious discussions of quantization are [Mackey, 1968] and [Souriau, 1969]; more recent and comprehensive treatments are, for example, [Woodhouse, 1992; Landsman, 1998], and [Ali and Englis, 2004].

4.1 Canonical quantization and systems of imprimitivity

The approach based on (2) is often called *canonical quantization*. Even apart from the issue of mathematical rigour, one can only side with Mackey [1992, 283], who wrote: ‘Simple and elegant as this model is, it appears at first sight to be quite arbitrary and ad hoc. It is difficult to understand how anyone could have guessed it and by no means obvious how to modify it to fit a model for space different from \mathbb{R}^r .’

One veil of the mystery of quantization was lifted by von Neumann [1931], who (following earlier heuristic proposals by Heisenberg, Schrödinger, Dirac, and Pauli) recognized that (2) does not merely provide a representation of the canonical commutation relations

$$(1) \quad [\mathcal{Q}_{\hbar}(p_j), \mathcal{Q}_{\hbar}(q^k)] = -i\hbar\delta_j^k,$$

but (subject to a regularity condition)⁹⁶ is *the only* such representation that is irreducible (up to unitary equivalence). In particular, the seemingly different formulations of quantum theory by Heisenberg and Schrödinger (amended by the inclusion of states and of observables, respectively — cf. Section 2) simply involved superficially different but unitarily equivalent representations of (1): the difference between matrices and waves was just one between coordinate systems in Hilbert space, so to speak. Moreover, any other conceivable formulation of quan-

⁹⁵The path integral approach to quantization is still under development and so far has had no impact on foundational debates, so we will not discuss it here. See [Albeverio and Høegh-Krohn, 1976] and [Glimm and Jaffe, 1987].

⁹⁶It is required that the unbounded operators $\mathcal{Q}_{\hbar}(p_j)$ and $\mathcal{Q}_{\hbar}(q^k)$ integrate to a unitary representation of the $2n + 1$ -dimensional Heisenberg group H_n , i.e. the unique connected and simply connected Lie group with $2n + 1$ -dimensional Lie algebra with generators X_i, Y_i, Z ($i = 1, \dots, n$) subject to the Lie brackets $[X_i, X_j] = [Y_i, Y_j] = 0$, $[X_i, Y_j] = \delta_{ij}Z$, $[X_i, Z] = [Y_i, Z] = 0$. Thus von Neumann’s uniqueness theorem for representations of the canonical commutation relations is (as he indeed recognized himself) really a uniqueness theorem for unitary representations of H_n for which the central element Z is mapped to $-i\hbar^{-1}1$, where $\hbar \neq 0$ is a *fixed* constant. See, for example, [Corwin and Greenleaf, 1989] or [Landsman, 1998].

tum mechanics — now simply *defined* as a (regular) Hilbert space representation of (1) — has to be equivalent to the one of Heisenberg and Schrödinger.⁹⁷

This, then, transfers the quantization problem of a particle moving on \mathbb{R}^n to the canonical commutation relations (1). Although a mathematically rigorous theory of these commutation relations (as they stand) exists [Jørgensen and Moore, 1984; Schmüdgen, 1990], they are problematic nonetheless. Firstly, technically speaking the operators involved are unbounded, and in order to represent physical observables they have to be self-adjoint; yet on their respective domains of self-adjointness the commutator on the left-hand side is undefined. Secondly, and more importantly, (1) relies on the possibility of choosing global coordinates on \mathbb{R}^n , which precludes a naive generalization to arbitrary configuration spaces. And thirdly, even if one has managed to quantize p and q by finding a representation of (1), the problem of quantizing other observables remains — think of the Hamiltonian and the Schrödinger equation.

About 50 years ago, Mackey set himself the task of making good sense of canonical quantization; see [Mackey, 1968; 1978; 1992] and the brief exposition below for the result. Although the author now regards Mackey's reformulation of quantization in terms of induced representations and systems of imprimitivity merely as a stepping stone towards our current understanding based on deformation theory and groupoids (cf. Subsection 4.3 below), Mackey's approach is (quite rightly) often used in the foundations of physics, and one is well advised to be familiar with it. In any case, Mackey [1992, 283] — continuing the previous quotation) claims with some justification that his approach to quantization 'removes much of the mystery.'

Like most approaches to quantization, Mackey assigns momentum and position a quite different role in quantum mechanics, despite the fact that in classical mechanics p and q can be interchanged by a canonical transformation.⁹⁸

1. The position operators $\mathcal{Q}_{\hbar}(q^j)$ are collectively replaced by a single projection-valued measure P on \mathbb{R}^n ,⁹⁹ which on $L^2(\mathbb{R}^n)$ is given by $P(E) = \chi_E$ as a multiplication operator. Given this P , any multiplication operator defined by a (measurable) function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ can be represented as $\int_{\mathbb{R}^n} dP(x) f(x)$, which is defined and self-adjoint on a suitable domain.¹⁰⁰ In particular, the position operators $\mathcal{Q}_{\hbar}(q^j)$ can be reconstructed from P by choosing $f(x) =$

⁹⁷This is unrelated to the issue of the Heisenberg picture versus the Schrödinger picture, which is about the time-evolution of observables versus that of states.

⁹⁸Up to a minus sign, that is. This is true globally on \mathbb{R}^n and locally on any symplectic manifold, where local Darboux coordinates do not distinguish between position and momentum.

⁹⁹A projection-valued measure P on a space Ω with Borel structure (i.e. equipped with a σ -algebra of measurable sets defined by the topology) with values in a Hilbert space \mathcal{H} is a map $E \mapsto P(E)$ from the Borel subsets $E \subset \Omega$ to the projections on \mathcal{H} that satisfies $P(\emptyset) = 0$, $P(\Omega) = 1$, $P(E)P(F) = P(F)P(E) = P(E \cap F)$ for all measurable $E, F \subset \Omega$, and $P(\cup_{i=1}^{\infty} E_i) = \sum_{i=1}^{\infty} P(E_i)$ for all countable collections of mutually disjoint $E_i \subset \Omega$.

¹⁰⁰This domain consists of all $\Psi \in \mathcal{H}$ for which $\int_{\mathbb{R}^n} d(\Psi, P(x)\Psi) |f(x)|^2 < \infty$.

x^j , i.e.

$$(2) \quad \mathcal{Q}_{\hbar}(q^j) = \int_{\mathbb{R}^n} dP(x) x^j.$$

2. The momentum operators $\mathcal{Q}_{\hbar}(p_j)$ are collectively replaced by a single unitary group representation $U(\mathbb{R}^n)$, defined on $L^2(\mathbb{R}^n)$ by

$$U(y)\Psi(x) := \Psi(x - y).$$

Each $\mathcal{Q}_{\hbar}(p_j)$ can be reconstructed from U by means of

$$(3) \quad \mathcal{Q}_{\hbar}(p_j)\Psi := i\hbar \lim_{t_j \rightarrow 0} t_j^{-1}(U(t_j) - 1)\Psi,$$

where $U(t_j)$ is U at $x^j = t_j$ and $x^k = 0$ for $k \neq j$.¹⁰¹

Consequently, it entails no loss of generality to work with the pair (P, U) instead of the pair $(\mathcal{Q}_{\hbar}(q^k), \mathcal{Q}_{\hbar}(p_j))$. The commutation relations (1) are now replaced by

$$(4) \quad U(x)P(E)U(x)^{-1} = P(xE),$$

where E is a (Borel) subset of \mathbb{R}^n and $xE = \{x\omega \mid \omega \in E\}$. On the basis of this reformulation, Mackey proposed the following sweeping generalization of the the canonical commutation relations:¹⁰²

A *system of imprimitivity* (\mathcal{H}, U, P) for a given action of a group G on a space Q consists of a Hilbert space \mathcal{H} , a unitary representation U of G on \mathcal{H} , and a projection-valued measure $E \mapsto P(E)$ on Q with values in \mathcal{H} , such that (4) holds for all $x \in G$ and all Borel sets $E \subset Q$.

In physics such a system describes the quantum mechanics of a particle moving on a configuration space Q on which G acts by symmetry transformations; see Subsection 4.3 for a more detailed discussion. When everything is smooth,¹⁰³ each element X of the Lie algebra \mathfrak{g} of G defines a generalized momentum operator

$$(5) \quad \mathcal{Q}_{\hbar}(X) = i\hbar dU(X)$$

on \mathcal{H} .¹⁰⁴ These operators satisfy the generalized canonical commutation relations¹⁰⁵

$$(6) \quad [\mathcal{Q}_{\hbar}(X), \mathcal{Q}_{\hbar}(Y)] = i\hbar \mathcal{Q}_{\hbar}([X, Y]).$$

¹⁰¹By Stone's theorem (cf. Reed and Simon, 1972), this operator is defined and self-adjoint on the set of all $\Psi \in H$ for which the limit exists.

¹⁰²All groups and spaces are supposed to be locally compact, and actions and representations are assumed continuous.

¹⁰³I.e. G is a Lie group, Q is a manifold, and the G -action is smooth.

¹⁰⁴This operator is defined and self-adjoint on the domain of vectors $\Psi \in \mathcal{H}$ for which $dU(X)\Psi := \lim_{t \rightarrow 0} t^{-1}(U(\exp(tX)) - 1)\Psi$ exists.

¹⁰⁵As noted before in the context of (1), the commutation relations (6), (8) and (9) do not hold on the domain of self-adjointness of the operators involved, but on a smaller common core.

Furthermore, in terms of the operators¹⁰⁶

$$(7) \quad \mathcal{Q}_\hbar(f) = \int_Q dP(x) f(x),$$

where f is a smooth function on Q and $X \in \mathfrak{g}$, one in addition has

$$(8) \quad [\mathcal{Q}_\hbar(X), \mathcal{Q}_\hbar(f)] = i\hbar \mathcal{Q}_\hbar(\xi_X^Q f),$$

where ξ_X^Q is the canonical vector field on Q defined by the G -action,¹⁰⁷ and

$$(9) \quad [\mathcal{Q}_\hbar(f_1), \mathcal{Q}_\hbar(f_2)] = 0.$$

Elementary quantum mechanics on \mathbb{R}^n corresponds to the special case $Q = \mathbb{R}^n$ and $G = \mathbb{R}^n$ with the usual additive group structure. To see this, we denote the standard basis of \mathbb{R}^3 (in its guise as the Lie algebra of \mathbb{R}^3) by the name (p_j) , and furthermore take $f_1(q) = q^j$, $f_2(q) = f(q) = q^k$. Eq. (6) for $X = p_j$ and $Y = p_k$ then reads $[\mathcal{Q}_\hbar(p_j), \mathcal{Q}_\hbar(p_k)] = 0$, eq. (8) yields the canonical commutation relations (1), and (9) states the commutativity of the position operators, i.e. $[\mathcal{Q}_\hbar(q^j), \mathcal{Q}_\hbar(q^k)] = 0$.

In order to incorporate spin, one picks $G = E(3) = SO(3) \times \mathbb{R}^3$ (i.e. the Euclidean motion group), acting on $Q = \mathbb{R}^3$ in the obvious (defining) way. The Lie algebra of $E(3)$ is $\mathbb{R}^6 = \mathbb{R}^3 \times \mathbb{R}^3$ as a vector space; we extend the basis (p_j) of the second copy of \mathbb{R}^3 (i.e. the Lie algebra of \mathbb{R}^3) by a basis (J_i) of the first copy of \mathbb{R}^3 (in its guise as the Lie algebra of $SO(3)$), and find that the $\mathcal{Q}_\hbar(J_i)$ are just the usual angular momentum operators.¹⁰⁸

Mackey’s generalization of von Neumann’s [1931] uniqueness theorem for the irreducible representations of the canonical commutation relations (1) is his *imprimitivity theorem*. This theorem applies to the special case where $Q = G/H$ for some (closed) subgroup $H \subset G$, and states that (up to unitary equivalence) there is a bijective correspondence between:

1. Systems of imprimitivity (\mathcal{H}, U, P) for the left-translation of G on G/H ;
2. Unitary representations U_χ of H .

This correspondence preserves irreducibility.¹⁰⁹

¹⁰⁶For the domain of $\mathcal{Q}_\hbar(f)$ see footnote 100.

¹⁰⁷I.e. $\xi_X^Q f(y) = d/dt|_{t=0}[f(\exp(-tX)y)]$.

¹⁰⁸The commutation relations in the previous paragraph are now extended by the familiar relations $[\mathcal{Q}_\hbar(J_i), \mathcal{Q}_\hbar(J_j)] = i\hbar \epsilon_{ijk} \mathcal{Q}_\hbar(J_k)$, $[\mathcal{Q}_\hbar(J_i), \mathcal{Q}_\hbar(p_j)] = i\hbar \epsilon_{ijk} \mathcal{Q}_\hbar(p_k)$, and $[\mathcal{Q}_\hbar(J_i), \mathcal{Q}_\hbar(q^j)] = i\hbar \epsilon_{ijk} \mathcal{Q}_\hbar(q^k)$.

¹⁰⁹Specifically, given U_χ the triple $(\mathcal{H}^\chi, U^\chi, P^\chi)$ is a system of imprimitivity, where $\mathcal{H}^\chi = L^2(G/H) \otimes \mathcal{H}_\chi$ carries the representation $U^\chi(G)$ induced by $U_\chi(H)$, and the P^χ act like multiplication operators. Conversely, if (\mathcal{H}, U, P) is a system of imprimitivity, then there exists a unitary representation $U_\chi(H)$ such that the triple (\mathcal{H}, U, P) is unitarily equivalent to the triple $(\mathcal{H}^\chi, U^\chi, P^\chi)$ just described. For example, for $G = E(3)$ and $H = SO(3)$ one has $\chi = j = 0, 1, 2, \dots$ and $\mathcal{H}^j = L^2(\mathbb{R}^3) \otimes \mathcal{H}_j$ (where $\mathcal{H}_j = \mathbb{C}^{2j+1}$ carries the given representation $U_j(SO(3))$).

For example, von Neumann's theorem is recovered as a special case of Mackey's by making the choice $G = \mathbb{R}^3$ and $H = \{e\}$ (so that $Q = \mathbb{R}^3$, as above): the uniqueness of the (regular) irreducible representation of the canonical commutation relations here follows from the uniqueness of the irreducible representation of the trivial group. A more illustrative example is $G = E(3)$ and $H = SO(3)$ (so that $Q = \mathbb{R}^3$), in which case the irreducible representations of the associated system of imprimitivity are classified by spin $j = 0, 1, \dots$. Mackey saw this as an explanation for the emergence of spin as a purely quantum-mechanical degree of freedom.¹¹⁰ Although the opinion that spin has no classical analogue was widely shared also among the pioneers of quantum theory,¹¹¹ it is now obsolete (see Subsection 4.3 below). Despite this unfortunate misinterpretation, Mackey's approach to canonical quantization is hard to surpass in power and clarity, and has many interesting applications.¹¹²

We mention one of specific interest to the philosophy of physics, namely the *Newton–Wigner position operator* (as analyzed by Wightman [1962]).¹¹³ Here the general question is whether a given unitary representation U of $G = E(3)$ on some Hilbert space \mathcal{H} may be extended to a system of imprimitivity with respect to $H = SO(3)$ (and hence $Q = \mathbb{R}^3$, as above); in that case, U (or rather the associated quantum system) is said to be *localizable* in \mathbb{R}^3 . Following Wigner's (1939) suggestion that a relativistic particle is described by an irreducible representation U of the Poincaré group P , one obtains a representation $U(E(3))$ by restricting $U(P)$ to the subgroup $E(3) \subset P$.¹¹⁴ It then follows from the previous analysis that the particle described by $U(P)$ is localizable if and only if $U(E(3))$ is induced by some representation of $SO(3)$. This can, of course, be settled, with the result that massive particles of arbitrary spin can be localized in \mathbb{R}^3 (the corresponding position operator being precisely the one of Newton and Wigner), whereas massless particles may be localized in \mathbb{R}^3 if and only if their helicity is less than one. In particular, the photon (and the graviton) cannot be localized in \mathbb{R}^3 in the stated

¹¹⁰By the usual arguments (Wigner's theorem), one may replace $SO(3)$ by $SU(2)$, so as to obtain $j = 0, 1/2, \dots$

¹¹¹This opinion goes back to Pauli [1925], who talked about a 'klassisch nicht beschreibbare Zweideutigkeit in den quantentheoretischen Eigenschaften des Elektrons,' (i.e. an 'ambivalence in the quantum theoretical properties of the electron that has no classical description') which was later identified as spin by Goudsmit and Uhlenbeck. Probably the first person to draw attention to the classical counterpart of spin was Souriau [1969]. Another misunderstanding about spin is that its ultimate explanation must be found in relativistic quantum mechanics.

¹¹²This begs the question about the 'best' possible proof of Mackey's imprimitivity theorem. Mackey's own proof was rather measure-theoretic in flavour, and did not shed much light on the origin of his result. Probably the shortest proof has been given by [Ørsted, 1979], but the insight brevity gives is still rather limited. Quite to the contrary, truly transparent proofs reduce a mathematical claim to a tautology. Such proofs, however, tend to require a formidable machinery to make this reduction work; see [Echterhoff *et al.*, 2002] and [Landsman, 2006a] for two different approaches to the imprimitivity theorem in this style.

¹¹³Fleming and Butterfield [2000] give an up-to-date introduction to particle localization in relativistic quantum theory. See also [De Bièvre, 2003].

¹¹⁴Strictly speaking, this hinges on the choice of an inertial frame in Minkowski space, with associated adapted co-ordinates such that the configuration space \mathbb{R}^3 in question is given by $x^0 = 0$.

sense.¹¹⁵

To appreciate our later material on both phase space quantization and deformation quantization, it is helpful to give a C^* -algebraic reformulation of Mackey’s approach. Firstly, by the spectral theorem [Reed and Simon, 1972; Pedersen, 1989], a projection-valued measure $E \mapsto P(E)$ on a space Q taking values in a Hilbert space \mathcal{H} is equivalent to a nondegenerate representation π of the commutative C^* -algebra $C_0(Q)$ on \mathcal{H} through the correspondence (7).¹¹⁶ Secondly, if \mathcal{H} in addition carries a unitary representation U of G , the defining condition (4) of a system of imprimitivity (given a G -action on Q) is equivalent to the covariance condition

$$(10) \quad U(x)\mathcal{Q}_{\hbar}(f)U(x)^{-1} = \mathcal{Q}_{\hbar}(L_x f)$$

for all $x \in G$ and $f \in C_0(Q)$, where $L_x f(m) = f(x^{-1}m)$. Thus a system of imprimitivity for a given G -action on Q is “the same” as a covariant nondegenerate representation of $C_0(Q)$. Thirdly, from a G -action on Q one can construct a certain C^* -algebra $C^*(G, Q)$, the so-called *transformation group C^* -algebra* defined by the action, which has the property that its nondegenerate representations correspond bijectively (and “naturally”) to covariant nondegenerate representations of $C_0(Q)$, and therefore to systems of imprimitivity for the given G -action [Effros and Hahn, 1967; Pedersen, 1979; Landsman, 1998]. In the C^* -algebraic approach to quantum physics, $C^*(G, Q)$ is the algebra of observables of a particle moving on Q subject to the symmetries defined by the G -action; its inequivalent irreducible representations correspond to the possible superselection sectors of the system [Doebner and Tolar, 1975; Majid, 1988; 1990; Landsman, 1990a; 1990b; 1992].¹¹⁷

4.2 Phase space quantization and coherent states

In Mackey’s approach to quantization, Q is the *configuration space* of the system; the associated position coordinates *commute* (cf. (9)). This is reflected by the correspondence just discussed between projection-valued measures on Q and representations of the *commutative* C^* -algebra $C_0(Q)$. The noncommutativity of observables (and the associated uncertainty relations) typical of quantum mechanics is incorporated by adding the symmetry group G to the picture and imposing

¹¹⁵Seeing photons as quantized light waves with two possible polarizations transverse to the direction of propagation, this last result is physically perfectly reasonable.

¹¹⁶A *representation* of a C^* -algebra \mathcal{A} on a Hilbert space \mathcal{H} is a linear map $\pi : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$ such that $\pi(AB) = \pi(A)\pi(B)$ and $\pi(A^*) = \pi(A)^*$ for all $A, B \in \mathcal{A}$. Such a representation is called *nondegenerate* when $\pi(A)\Psi = 0$ for all $A \in \mathcal{A}$ implies $\Psi = 0$.

¹¹⁷Another reformulation of Mackey’s approach, or rather an extension of it, has been given by Isham [1984]. In an attempt to reduce the whole theory to a problem in group representations, he proposed that the possible quantizations of a particle with configuration space G/H are given by the inequivalent irreducible representations of a “canonical group” $G_c = G \ltimes V$, where V is the lowest-dimensional vector space that carries a representation of G under which G/H is an orbit in the dual vector space V^* . All pertinent systems of imprimitivity then indeed correspond to unitary representations of G_c , but this group has many other representations whose physical interpretation is obscure. See also footnote 157.

the relations (4) (or, equivalently, (8) or (10)). As we have pointed out, this procedure upsets the symmetry between the phase space variables position and momentum in classical mechanics.

This somewhat unsatisfactory feature of Mackey's approach may be avoided by replacing Q by the *phase space* of the system, henceforth called M .¹¹⁸ In this approach, noncommutativity is incorporated by a treacherously tiny modification to Mackey's setup. Namely, the projection-valued measure $E \mapsto P(M)$ on M with which he starts is now replaced by a *positive-operator-valued measure* or *POVM* on M , still taking values in some Hilbert space \mathcal{K} . This is a map $E \mapsto A(E)$ from the (Borel) subsets E of M to the collection of *positive* bounded operators on \mathcal{K} ,¹¹⁹ satisfying $A(\emptyset) = 0$, $A(M) = 1$, and $A(\cup_i E_i) = \sum_i A(E_i)$ for any countable collection of disjoint Borel sets E_i .¹²⁰ A POVM that satisfies $A(E \cap F) = A(E)A(F)$ for all (Borel) $E, F \subset M$ is precisely a projection-valued measure, so that a POVM is a generalization of the latter.¹²¹ The point, then, is that *a given POVM defines a quantization procedure* by the stipulation that a classical observable f (i.e. a measurable function on the phase space M , for simplicity assumed bounded) is quantized by the operator¹²²

$$(11) \quad \mathcal{Q}(f) = \int_M dA(x)f(x).$$

Thus the seemingly slight move from projection-valued measures on configuration space to positive-operator valued measures on phase space gives a wholly new perspective on quantization, actually reducing this task to the problem of finding such POVM's.¹²³

¹¹⁸Here the reader may think of the simplest case $M = \mathbb{R}^6$, the space of p 's and q 's of a particle moving on \mathbb{R}^3 . More generally, if Q is the configuration space, the associated phase space is the cotangent bundle $M = T^*Q$. Even more general phase spaces, namely arbitrary symplectic manifolds, may be included in the theory as well. References for what follows include [Busch *et al.*, 1998; Schroeck, 1996], and [Landsman, 1998; 1999a].

¹¹⁹A bounded operator A on \mathcal{K} is called *positive* when $(\Psi, A\Psi) \geq 0$ for all $\Psi \in \mathcal{K}$. Consequently, it is self-adjoint with spectrum contained in \mathbb{R}^+ .

¹²⁰Here the infinite sum is taken in the weak operator topology. Note that the above conditions force $0 \leq A(E) \leq 1$, in the sense that $0 \leq (\Psi, A(E)\Psi) \leq (\Psi, \Psi)$ for all $\Psi \in \mathcal{K}$.

¹²¹This has given rise to the so-called *operational approach* to quantum theory, in which observables are not represented by self-adjoint operators (or, equivalently, by their associated projection-valued measures), but by POVM's. The space M on which the POVM is defined is the space of outcomes of the measuring instrument; the POVM is determined by both A and a calibration procedure for this instrument. The probability that in a state ρ the outcome of the experiment lies in $E \subset M$ is taken to be $\text{Tr}(\rho A(E))$. See [Davies, 1976; Holevo, 1982; Ludwig, 1985; Schroeck, 1996; Busch *et al.*, 1998], and [De Muynck, 2002].

¹²²The easiest way to define the right-hand side of (11) is to fix $\Psi \in \mathcal{K}$ and define a probability measure p_Ψ on M by means of $p_\Psi(E) = (\Psi, A(E)\Psi)$. One then *defines* $\mathcal{Q}(f)$ as an operator through its expectation values $(\Psi, \mathcal{Q}(f)\Psi) = \int_M dp_\Psi(x)f(x)$. The expression (11) generalizes (7), and also generalizes the spectral resolution of the operator $f(A) = \int_{\mathbb{R}} dP(\lambda)f(\lambda)$, where P is the projection-valued measure defined by a self-adjoint operator A .

¹²³An important feature of \mathcal{Q} is that it is *positive* in the sense that if $f(x) \geq 0$ for all $x \in M$, then $(\Psi, \mathcal{Q}(f)\Psi) \geq 0$ for all $\Psi \in \mathcal{K}$. In other words, \mathcal{Q} is positive as a map from the C^* -algebra $C_0(M)$ to the C^* -algebra $\mathcal{B}(\mathcal{H})$.

The solution to this problem is greatly facilitated by *Naimark's dilation theorem*.¹²⁴ This states that, given a POVM $E \mapsto A(E)$ on M in a Hilbert space \mathcal{K} , there exists a Hilbert space \mathcal{H} carrying a projection-valued measure P on M and an isometric injection $\mathcal{K} \hookrightarrow \mathcal{H}$, such that

$$(12) \quad A(E) = [\mathcal{K}]P(E)[\mathcal{K}]$$

for all $E \subset M$ (where $[\mathcal{K}]$ is the orthogonal projection from \mathcal{H} onto \mathcal{K}). Combining this with Mackey's imprimitivity theorem yields a powerful generalization of the latter [Poulsen, 1970; Neumann, 1972; Scutaru, 1977; Cattaneo, 1979; Castrigiano and Henrichs, 1980]. First, define a *generalized system of imprimitivity* (\mathcal{K}, U, A) for a given action of a group G on a space M as a POVM A on M taking values in a Hilbert space \mathcal{K} , along with a unitary representation V of G on \mathcal{K} such that

$$(13) \quad V(x)A(E)V(x)^{-1} = A(xE)$$

for all $x \in G$ and $E \subset M$; cf. (4). Now assume $M = G/H$ (and the associated canonical left-action on M). The *generalized imprimitivity theorem* states that a generalized system of imprimitivity (\mathcal{K}, V, A) for this action is necessarily (unitarily equivalent to) a reduction of a system of imprimitivity (\mathcal{H}, U, P) for the same action. In other words, the Hilbert space \mathcal{H} in Naimark's theorem carries a unitary representation $U(G)$ that commutes with the projection $[\mathcal{K}]$, and the representation $V(G)$ is simply the restriction of U to \mathcal{K} . Furthermore, the POVM A has the form (12). The structure of (\mathcal{H}, U, P) is fully described by Mackey's imprimitivity theorem, so that one has a complete classification of generalized systems of imprimitivity.¹²⁵ One has

$$(14) \quad \mathcal{K} = p\mathcal{H}; \quad \mathcal{H} = L^2(M) \otimes \mathcal{H}_\chi,$$

where L^2 is defined with respect to a suitable measure on $M = G/H$,¹²⁶ the Hilbert space \mathcal{H}_χ carries a unitary representation of H , and p is a projection in the commutant of the representation $U^\times(G)$ induced by $U_\chi(G)$.¹²⁷ The quantization (11) is given by

$$(15) \quad \mathcal{Q}(f) = pfp,$$

where f acts on $L^2(M) \otimes \mathcal{H}_\chi$ as a multiplication operator, i.e. $(f\Psi)(x) = f(x)\Psi(x)$. In particular, one has $P(E) = \chi_E$ (as a multiplication operator) for a region $E \subset M$ of phase space, so that $\mathcal{Q}(\chi_E) = A(E)$. Consequently, the probability

¹²⁴See, for example, [Riesz and Sz.-Nagy, 1990]. It is better, however, to see Naimark's theorem as a special case of Stinesprings's, as explained e.g. in [Landsman, 1998], and below.

¹²⁵Continuing footnote 109: $V(G)$ is necessarily a subrepresentation of some representation $U^\times(G)$ induced by $U_\chi(H)$.

¹²⁶In the physically relevant case that G/H is symplectic (so that it typically is a coadjoint orbit for G) one should take a multiple of the Liouville measure.

¹²⁷The explicit form of $U^\times(g)$, $g \in G$, depends on the choice of a cross-section $\sigma : G/H \rightarrow G$ of the projection $\pi : G \rightarrow G/H$ (i.e. $\pi \circ \sigma = \text{id}$). If the measure on G/H defining $L^2(G/H)$ is G -invariant, the explicit formula is $U^\times(g)\Psi(x) = U_\chi(s(x)^{-1}gs(g^{-1}x))\Psi(g^{-1}x)$.

that in a state ρ (i.e. a density matrix on \mathcal{K}) the system is localized in E is given by $\text{Tr}(\rho A(E))$.

In a more natural way than in Mackey’s approach, the covariant POVM quantization method allows one to incorporate space-time symmetries *ab initio* by taking G to be the Galilei group or the Poincaré group, and choosing H such that G/H is a physical phase space (on which G , then, canonically acts). See [Ali *et al.*, 1995] and [Schroeck, 1996].

Another powerful method of constructing POVM’s on phase space (which in the presence of symmetries overlaps with the preceding one)¹²⁸ is based on *coherent states*.¹²⁹ The minimal definition of coherent states in a Hilbert space \mathcal{H} for a phase space M is that (for some fixed value of Planck’s constant \hbar , for the moment) one has an injection¹³⁰ $M \hookrightarrow \mathcal{H}$, $z \mapsto \Psi_z^\hbar$, such that

$$(16) \quad \|\Psi_z^\hbar\| = 1$$

for all $z \in M$, and

$$(17) \quad c_\hbar \int_M d\mu_L(z) |(\Psi_z^\hbar, \Phi)|^2 = 1,$$

for each $\Phi \in \mathcal{H}$ of unit norm (here μ_L is the Liouville measure on M and $c_\hbar > 0$ is a suitable constant).¹³¹ Condition (17) guarantees that we may define a POVM on M in \mathcal{K} by¹³²

$$(18) \quad A(E) = c_\hbar \int_E d\mu_L(z) [\Psi_z^\hbar].$$

Eq. (11) then simply reads (inserting the \hbar -dependence of \mathcal{Q} and a suffix B for later use)

$$(19) \quad \mathcal{Q}_\hbar^B(f) = c_\hbar \int_M d\mu_L(z) f(z) [\Psi_z^\hbar].$$

The time-honoured example, due to Schrödinger [1926b], is $M = \mathbb{R}^{2n}$, $\mathcal{H} = L^2(\mathbb{R}^n)$, and

$$(20) \quad \Psi_{(p,q)}^\hbar(x) = (\pi\hbar)^{-n/4} e^{-ipq/2\hbar} e^{ipx/\hbar} e^{-(x-q)^2/2\hbar}.$$

¹²⁸Suppose there is a vector $\Omega \in \mathcal{K}$ such that $\int_{G/H} d\mu(x) |(\Omega, V(\sigma(x))\Omega)|^2 < \infty$ with respect to some cross-section $\sigma : G/H \rightarrow G$ and a G -invariant measure μ , as well as $V(h)\Omega = U_\chi(h)\Omega$ for all $h \in H$, where $U_\chi : H \rightarrow \mathbb{C}$ is *one-dimensional*. Then (taking $\hbar = 1$) the vectors $V(\sigma(x))\Omega$ (suitably normalized) form a family of coherent states on G/H [Ali *et al.*, 1995; Schroeck, 1996; Ali, Antoine, and Gazeau, 2000]. For example, the coherent states (20) are of this form for the Heisenberg group.

¹²⁹ See [Klauder and Skagerstam, 1985; Perelomov, 1986; Odziejewicz, 1992; Paul and Uribe, 1995; 1996; Ali *et al.*, 1995], and [Ali *et al.*, 2000], for general discussions of coherent states.

¹³⁰This injection must be continuous as a map from M to $\mathbb{P}\mathcal{H}$, the projective Hilbert space of \mathcal{H} .

¹³¹Other measures might occur here; see, for example, [Bonechi and De Bièvre, 2000].

¹³²Recall that $[\Psi]$ is the orthogonal projection onto a unit vector Ψ .

Eq. (17) then holds with $d\mu_L(p, q) = (2\pi)^{-n}d^npd^nq$ and $c_{\hbar} = \hbar^{-n}$. One may verify that $\mathcal{Q}_{\hbar}^B(p_j)$ and $\mathcal{Q}_{\hbar}^B(q^j)$ coincide with Schrödinger’s operators (2). This example illustrates that coherent states need *not* be mutually orthogonal; in fact, in terms of $z = p + iq$ one has for the states in (20)

$$(21) \quad |(\Psi_z^{\hbar}, \Psi_w^{\hbar})|^2 = e^{-|z-w|^2/2\hbar},$$

the significance of this result will emerge later on.

In the general case, it is an easy matter to verify Naimark’s dilation theorem for the POVM (18): changing notation so that the vectors Ψ_z^{\hbar} now lie in \mathcal{K} , one finds

$$(22) \quad \mathcal{H} = L^2(M, c_{\hbar}\mu_L),$$

the embedding $W : \mathcal{K} \hookrightarrow \mathcal{H}$ being given by $(W\Phi)(z) = (\Psi_z^{\hbar}, \Phi)$. The projection-valued measure P on \mathcal{H} is just $P(E) = \chi_E$ (as a multiplication operator), and the projection p onto $W\mathcal{K}$ is given by

$$(23) \quad p\Psi(z) = c_{\hbar} \int_M d\mu_L(w) (\Psi_z^{\hbar}, \Psi_w^{\hbar}) \Psi(w).$$

Consequently, (19) is unitarily equivalent to (15), where f acts on $L^2(M)$ as a multiplication operator.¹³³

*Thus (15) and (22) (or its possible extension (14)) form the essence of phase space quantization.*¹³⁴

We close this subsection in the same fashion as the previous one, namely by pointing out the C^* -algebraic significance of POVM’s. This is extremely easy: whereas a projection-valued measure on M in \mathcal{H} is the same as a nondegenerate representation of $C_0(M)$ in \mathcal{H} , a POVM on M in a Hilbert space \mathcal{K} is nothing but a nondegenerate *completely positive map* $\varphi : C_0(M) \rightarrow \mathcal{B}(\mathcal{K})$.¹³⁵ Consequently, Naimark’s dilation theorem becomes a special case of Stinespring’s [1955] theorem: if $\mathcal{Q} : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{K})$ is a completely positive map, there exists a Hilbert space \mathcal{H} carrying a representation π of $C_0(M)$ and an isometric injection $\mathcal{K} \hookrightarrow \mathcal{H}$, such that $\mathcal{Q}(f) = [\mathcal{K}]\pi(f)[\mathcal{K}]$ for all $f \in C_0(M)$. In terms of $\mathcal{Q}(C_0(M))$, the covariance condition (13) becomes $U(x)\mathcal{Q}(f)U(x)^{-1} = \mathcal{Q}(L_x f)$, just like (10).

¹³³This leads to a close relationship between coherent states and Hilbert spaces with a reproducing kernel; see [Landsman, 1998] or [Ali *et al.*, 2000].

¹³⁴See also footnote 172 below.

¹³⁵A map $\varphi : \mathcal{A} \rightarrow \mathcal{B}$ between C^* -algebras is called positive when $\varphi(A) \geq 0$ whenever $A \geq 0$; such a map is called *completely positive* if for all $n \in \mathbb{N}$ the map $\varphi_n : \mathcal{A} \otimes M_n(\mathbb{C}) \rightarrow \mathcal{B} \otimes M_n(\mathbb{C})$, defined by linear extension of $\varphi \otimes \text{id}$ on elementary tensors, is positive (here $M_n(\mathbb{C})$ is the C^* -algebra of $n \times n$ complex matrices). When \mathcal{A} is commutative a nondegenerate positive map $\mathcal{A} \rightarrow \mathcal{B}$ is automatically completely positive for any \mathcal{B} .

4.3 Deformation quantization

So far, we have used the word ‘quantization’ in a heuristic way, basing our account on historical continuity rather than on axiomatic foundations. In this subsection and the next we set the record straight by introducing two alternative ways of looking at quantization in an axiomatic way. We start with the approach that historically came last, but which conceptually is closer to the material just discussed. This is *deformation quantization*, originating in the work of Berezin [1974; 1975a; 1975b], Vey [1975], and Bayen *et al.* [1977]. We here follow the C^* -algebraic approach to deformation quantization proposed by Rieffel [1989a; 1994], since it is not only mathematically transparent and rigorous, but also reasonably close to physical practice.¹³⁶ Due to the mathematical language used, this method of course naturally fits into the general C^* -algebraic approach to quantum physics.

The key idea of deformation quantization is that quantization should be defined through the property of having the correct classical limit. Consequently, Planck’s “constant” \hbar is treated as a variable, so that for each of its values one should have a quantum theory. The key requirement is that this family of quantum theories converges to the underlying classical theory as $\hbar \rightarrow 0$.¹³⁷ The mathematical implementation of this idea is quite beautiful, in that the classical algebra of observables is “glued” to the family of quantum algebras of observables in such a way that the classical theory literally forms the boundary of the space containing the pertinent quantum theories (one for each value of $\hbar > 0$). Technically, this is done through the concept of a *continuous field of C^* -algebras*.¹³⁸ What follows may sound unnecessarily technical, but the last 15 years have indicated that this yields exactly the right definition of quantization.

Let $I \subset \mathbb{R}$ be the set in which \hbar takes values; one usually has $I = [0, 1]$, but when the phase space is compact, \hbar often takes values in a countable subset of $(0, 1]$.¹³⁹ The same situation occurs in the theory of infinite systems; see Section 6. In any case, I should contain zero as an accumulation point. A continuous field of C^* -algebras over I , then, consists of a C^* -algebra \mathcal{A} , a collection of C^* -algebras $\{\mathcal{A}_{\hbar}\}_{\hbar \in I}$, and a surjective morphism $\varphi_{\hbar} : \mathcal{A} \rightarrow \mathcal{A}_{\hbar}$ for each $\hbar \in I$, such that:

1. The function $\hbar \mapsto \|\varphi_{\hbar}(A)\|_{\hbar}$ is in $C_0(I)$ for all $A \in \mathcal{A}$;¹⁴⁰
2. The norm of any $A \in \mathcal{A}$ is $\|A\| = \sup_{\hbar \in I} \|\varphi_{\hbar}(A)\|$;

¹³⁶See also [Landsman, 1998] for an extensive discussion of the C^* -algebraic approach to deformation quantization. In other approaches to deformation quantization, such as the theory of star products, \hbar is a formal parameter rather than a real number. In particular, the meaning of the limit $\hbar \rightarrow 0$ is obscure.

¹³⁷Cf. the preamble to Section 5 for further comments on this limit.

¹³⁸See [Dixmier, 1977; Fell and Doran, 1988], and [Kirchberg and Wassermann, 1995] for three different approaches to the same concept. Our definition follows the latter; replacing I by an arbitrary locally compact Hausdorff space one finds the general definition.

¹³⁹Cf. [Landsman, 1998] and footnote 204, but in contrast see [Rieffel, 1989a] for the example of the noncommutative torus, where one quantizes a compact phase space for each $\hbar \in (0, 1]$. Further examples of this sort are discussed by Natsume and Nest [1999], Natsume, Nest and Ingo [2003], and Hawkins [2005].

¹⁴⁰Here $\|\cdot\|_{\hbar}$ is the norm in the C^* -algebra \mathcal{A}_{\hbar} .

3. For any $f \in C_0(I)$ and $A \in \mathcal{A}$ there is an element $fA \in \mathcal{A}$ for which $\varphi_{\hbar}(fA) = f(\hbar)\varphi_{\hbar}(A)$ for all $\hbar \in I$.

The idea is that the family $(\mathcal{A}_{\hbar})_{\hbar \in I}$ of C^* -algebras is glued together by specifying a topology on the bundle $\coprod_{\hbar \in [0,1]} \mathcal{A}_{\hbar}$ (disjoint union). However, this topology is in fact defined rather indirectly, via the specification of the space of continuous sections of the bundle.¹⁴¹ Namely, a *continuous section* of the field is by definition an element $\{A_{\hbar}\}_{\hbar \in I}$ of $\prod_{\hbar \in I} \mathcal{A}_{\hbar}$ (equivalently, a map $\hbar \mapsto A_{\hbar}$ where $A_{\hbar} \in \mathcal{A}_{\hbar}$) for which there is an $A \in \mathcal{A}$ such that $A_{\hbar} = \varphi_{\hbar}(A)$ for all $\hbar \in I$. It follows that the C^* -algebra \mathcal{A} may actually be identified with the space of continuous sections of the field: if we do so, the morphism φ_{\hbar} is just the evaluation map at \hbar .¹⁴²

Physically, \mathcal{A}_0 is the commutative algebra of observables of the underlying classical system, and for each $\hbar > 0$ the noncommutative C^* -algebra \mathcal{A}_{\hbar} is supposed to be the algebra of observables of the corresponding quantum system at value \hbar of Planck’s constant. The algebra \mathcal{A}_0 , then, is of the form $C_0(M)$, where M is the phase space defining the classical theory. A phase space has more structure than an arbitrary topological space; it is a manifold on which a Poisson bracket $\{, \}$ can be defined. For example, on $M = \mathbb{R}^{2n}$ one has the familiar expression

$$(24) \quad \{f, g\} = \sum_j \frac{\partial f}{\partial p_j} \frac{\partial g}{\partial q^j} - \frac{\partial f}{\partial q^j} \frac{\partial g}{\partial p_j}.$$

Technically, M is taken to be a *Poisson manifold*. This is a manifold equipped with a Lie bracket $\{, \}$ on $C^\infty(M)$ with the property that for each $f \in C^\infty(M)$ the map $g \mapsto \{f, g\}$ defines a derivation of the commutative algebra structure of $C^\infty(M)$ given by pointwise multiplication. Hence this map is given by a vector field ξ_f , called the *Hamiltonian vector field* of f (i.e. one has $\xi_f g = \{f, g\}$). *Symplectic manifolds* are special instances of Poisson manifolds, characterized by the property that the Hamiltonian vector fields exhaust the tangent bundle. A Poisson manifold is foliated by its *symplectic leaves*: a given symplectic leaf L is characterized by the property that at each $x \in L$ the tangent space $T_x L \subset T_x M$ is spanned by the collection of all Hamiltonian vector fields at x . Consequently, the flow of any Hamiltonian vector field on M through a given point lies in its entirety within the symplectic leaf containing that point. The simplest example of a Poisson manifold is $M = \mathbb{R}^{2n}$ with Poisson bracket (24); this manifold is even symplectic.¹⁴³

¹⁴¹This is reminiscent of the Gelfand–Naimark theorem for commutative C^* -algebras, which specifies the topology on a locally compact Hausdorff space X via the C^* -algebra $C_0(X)$. Similarly, in the theory of (locally trivial) vector bundles the Serre–Swan theorem allows one to reconstruct the topology on a vector bundle $E \xrightarrow{\pi} X$ from the space $\Gamma_0(E)$ of continuous sections of E , seen as a (finitely generated projective) $C_0(X)$ -module. See, for example, [Gracia-Bondía *et al.*, 2001]. The third condition in our definition of a continuous field of C^* -algebras makes \mathcal{A} a $C_0(I)$ -module in the precise sense that there exists a nondegenerate morphism from $C_0(I)$ to the center of the multiplier of \mathcal{A} . This property may also replace our condition 3.

¹⁴²The structure of \mathcal{A} as a C^* -algebra corresponds to the operations of pointwise scalar multiplication, addition, adjointing, and operator multiplication on sections.

¹⁴³See [Marsden and Ratiu, 1994] for a mechanics-oriented introduction to Poisson manifolds; also cf. [Landsman, 1998] or [Butterfield, 2005] for the basic facts. A classical mathematical paper on Poisson manifolds is [Weinstein, 1983].

After this preparation, our basic definition is this:¹⁴⁴

A deformation quantization of a phase space M consists of a continuous field of C^ -algebras $(\mathcal{A}_{\hbar})_{\hbar \in [0,1]}$ (with $\mathcal{A}_0 = C_0(M)$), along with a family of self-adjoint¹⁴⁵ linear maps $\mathcal{Q}_{\hbar} : C_c^{\infty}(M) \rightarrow \mathcal{A}_{\hbar}$, $\hbar \in (0, 1]$, such that:*

1. For each $f \in C_c^{\infty}(M)$ the map defined by $0 \mapsto f$ and $\hbar \mapsto \mathcal{Q}_{\hbar}(f)$ ($\hbar \neq 0$) is a continuous section of the given continuous field;¹⁴⁶
2. For all $f, g \in C_c^{\infty}(M)$ one has

$$(25) \quad \lim_{\hbar \rightarrow 0} \left\| \frac{i}{\hbar} [\mathcal{Q}_{\hbar}(f), \mathcal{Q}_{\hbar}(g)] - \mathcal{Q}_{\hbar}(\{f, g\}) \right\|_{\hbar} = 0.$$

Obvious continuity properties one might like to impose, such as

$$(26) \quad \lim_{\hbar \rightarrow 0} \|\mathcal{Q}_{\hbar}(f)\mathcal{Q}_{\hbar}(g) - \mathcal{Q}_{\hbar}(fg)\| = 0,$$

or

$$(27) \quad \lim_{\hbar \rightarrow 0} \|\mathcal{Q}_{\hbar}(f)\| = \|f\|_{\infty},$$

turn out to be an automatic consequence of this definition.¹⁴⁷ Condition (25), however, transcends the C^* -algebraic setting, and is the key ingredient in proving (among other things) that the quantum dynamics converges to the classical dynamics;¹⁴⁸ see Section 5. The map \mathcal{Q}_{\hbar} is the quantization map at value \hbar of Planck's constant; we feel it is the most precise formulation of Heisenberg's original *Umdeutung* of classical observables known to date. It has the same interpretation as the heuristic symbol \mathcal{Q}_{\hbar} used so far: the operator $\mathcal{Q}_{\hbar}(f)$ is the quantum-mechanical observable whose classical counterpart is f .

This has turned out to be an fruitful definition of quantization, firstly because most well-understood examples of quantization fit into it [Rieffel, 1994; Landsman, 1998], and secondly because it has suggested various fascinating new ones [Rieffel,

¹⁴⁴Here $C_c^{\infty}(M)$ stands for the space of smooth functions on M with compact support; this is a norm-dense subalgebra of $\mathcal{A}_0 = C_0(M)$. The question whether the maps \mathcal{Q}_{\hbar} can be extended from $C_c^{\infty}(M)$ to $C_0(M)$ has to be answered on a case by case basis. Upon such an extension, if it exists, condition (25) will lose its meaning, since the Poisson bracket $\{f, g\}$ is not defined for all $f, g \in C_0(M)$.

¹⁴⁵I.e. $\mathcal{Q}_{\hbar}(\bar{f}) = \mathcal{Q}_{\hbar}(f)^*$.

¹⁴⁶Equivalently, one could extend the family $(\mathcal{Q}_{\hbar})_{\hbar \in (0,1]}$ to $\hbar = 0$ by $\mathcal{Q}_0 = \text{id}$, and state that $\hbar \mapsto \mathcal{Q}_{\hbar}(f)$ is a continuous section. Also, one could replace this family of maps by a single section $\mathcal{Q} : C_c^{\infty}(M) \rightarrow \mathcal{A}$ of φ_0 and define $\mathcal{Q}_{\hbar} = \varphi_{\hbar} \circ \mathcal{Q} : C_c^{\infty}(M) \rightarrow \mathcal{A}_{\hbar}$.

¹⁴⁷That they are automatic should not distract from the fact that especially (27) is a beautiful connection between classical and quantum mechanics. See footnote 89 for the meaning of $\|f\|_{\infty}$.

¹⁴⁸This insight is often attributed to Dirac [1930], who was the first to recognize the analogy between the commutator in quantum mechanics and the Poisson bracket in classical mechanics. In fact, the Poisson structure on M is uniquely determined by the continuous field structure together with condition (25). Thus the choice of the \mathcal{Q}_{\hbar} is secondary.

1989a; Natsume and Nest, 1999; Natsume *et al.*, 2003; Hawkins, 2005]. Restricting ourselves to the former, we note, for example, that (19) with (20) defines a deformation quantization of the phase space \mathbb{R}^{2n} (with standard Poisson bracket) if one takes \mathcal{A}_\hbar to be the C^* -algebra of compact operators on the Hilbert space $L^2(\mathbb{R}^n)$. This is called the *Berezin quantization* of \mathbb{R}^{2n} (as a phase space),¹⁴⁹ explicitly, for $\Phi \in L^2(\mathbb{R}^n)$ one has

$$(28) \quad \mathcal{Q}_\hbar^B(f)\Phi(x) = \int_{\mathbb{R}^{2n}} \frac{d^n p d^n q d^n y}{(2\pi\hbar)^n} f(p, q) \overline{\Psi_{(p,q)}^\hbar(y)} \Phi(y) \Psi_{(p,q)}^\hbar(x).$$

This quantization has the distinguishing feature of positivity,¹⁵⁰ a property not shared by its more famous sister called *Weyl quantization*.¹⁵¹ The latter is a deformation quantization of \mathbb{R}^{2n} as well, having the same continuous field of C^* -algebras, but differing from Berezin quantization in its quantization map

$$(29) \quad \mathcal{Q}_\hbar^W(f)\Phi(x) = \int_{\mathbb{R}^{2n}} \frac{d^n p d^n q}{(2\pi\hbar)^n} e^{ip(x-q)/\hbar} f(p, \frac{1}{2}(x+q)) \Phi(q).$$

Although it lacks good positivity and hence continuity properties,¹⁵² Weyl quantization enjoys better symmetry properties than Berezin quantization.¹⁵³ Despite these differences, which illustrate the lack of uniqueness of concrete quantization procedures, Weyl and Berezin quantization both reproduce Schrödinger’s position and momentum operators (2).¹⁵⁴ Furthermore, if $f \in L^1(\mathbb{R}^{2n})$, then $\mathcal{Q}_\hbar^B(f)$ and $\mathcal{Q}_\hbar^W(f)$ are trace class, with

$$(30) \quad \text{Tr } \mathcal{Q}_\hbar^B(f) = \text{Tr } \mathcal{Q}_\hbar^W(f) = \int_{\mathbb{R}^{2n}} \frac{d^n p d^n q}{(2\pi\hbar)^n} f(p, q).$$

Weyl and Berezin quantization are related by

$$(31) \quad \mathcal{Q}_\hbar^B(f) = \mathcal{Q}_\hbar^W(e^{\frac{\hbar}{4}\Delta_{2n}} f),$$

¹⁴⁹In the literature, Berezin quantization on \mathbb{R}^{2n} is often called anti-Wick quantization (or ordering), whereas on compact complex manifolds it is sometimes called Toeplitz or Berezin–Toeplitz quantization. Coherent states based on other phase spaces often define deformation quantizations as well; see [Landsman, 1998].

¹⁵⁰Cf. footnote 123. As a consequence, (28) is valid not only for $f \in C_c^\infty(\mathbb{R}^{2n})$, but even for all $f \in L^\infty(\mathbb{R}^{2n})$, and the extension of \mathcal{Q}_\hbar^B from $C_c^\infty(\mathbb{R}^{2n})$ to $L^\infty(\mathbb{R}^{2n})$ is continuous.

¹⁵¹The original reference is Weyl [1931]. See, for example, [Dubin *et al.*, 2000] and [Esposito *et al.*, 2004] for a modern physics-oriented yet mathematically rigorous treatment. See also [Rieffel, 1994] and [Landsman, 1998] for a discussion from the perspective of deformation quantization, as well as [Binz *et al.*, 2004] for infinite-dimensional examples.

¹⁵²Nonetheless, Weyl quantization may be extended from $C_c^\infty(\mathbb{R}^{2n})$ to much larger function spaces using techniques from the theory of distributions (leaving the Hilbert space setting typical of quantum mechanics). The classical treatment is in Hörmander [1979; 1985a].

¹⁵³Weyl quantization is covariant under the affine symplectic group $\text{Sp}(n, \mathbb{R}) \ltimes \mathbb{R}^{2n}$, whereas Berezin quantization is merely covariant under its subgroup $\text{O}(2n) \ltimes \mathbb{R}^{2n}$.

¹⁵⁴This requires a formal extension of the maps \mathcal{Q}_\hbar^W and \mathcal{Q}_\hbar^B to unbounded functions on M like p_j and q^j .

where $\Delta_{2n} = \sum_{j=1}^n (\partial^2/\partial p_j^2 + \partial^2/\partial(q^j)^2)$, from which it may be shown that Weyl and Berezin quantization are *asymptotically equal* in the sense that for any $f \in C_c^\infty(\mathbb{R}^{2n})$ one has

$$(32) \quad \lim_{\hbar \rightarrow 0} \|\mathcal{Q}_\hbar^B(f) - \mathcal{Q}_\hbar^W(f)\| = 0.$$

Mackey’s approach to quantization also finds its natural home in the setting of deformation quantization. Let a Lie group G act on a manifold Q , interpreted as a *configuration space*, as in Subsection 4.1. It turns out that the corresponding classical *phase space* is the manifold $\mathfrak{g}^* \times Q$, equipped with the so-called *semidirect product Poisson structure* [Marsden *et al.*, 1984; Krishnaprasad and Marsden, 1987]. Relative to a basis (T_a) of the Lie algebra \mathfrak{g} of G with structure constants C_{ab}^c (i.e. $[T_a, T_b] = \sum_c C_{ab}^c T_c$), the Poisson bracket in question is given by

$$(33) \quad \{f, g\} = \sum_a \left(\xi_a^M f \frac{\partial g}{\partial \theta_a} - \frac{\partial f}{\partial \theta_a} \xi_a^M g \right) - \sum_{a,b,c} C_{ab}^c \theta_c \frac{\partial f}{\partial \theta_a} \frac{\partial g}{\partial \theta_b},$$

where $\xi_a^M = \xi_{T_a}^M$. To illustrate the meaning of this lengthy expression, we consider a few special cases. First, take $f = X \in \mathfrak{g}$ and $g = Y \in \mathfrak{g}$ (seen as linear functions on the dual \mathfrak{g}^*). This yields

$$(34) \quad \{X, Y\} = -[X, Y].$$

Subsequently, assume that g depends on position q alone. This leads to

$$(35) \quad \{X, g\} = -\xi_X^M g.$$

Finally, assume that $f = f_1$ and $g = f_2$ depend on q only; this clearly gives

$$(36) \quad \{f_1, f_2\} = 0.$$

The two simplest physically relevant examples, already considered at the quantum level in Subsection 4.1, are as follows. First, take $G = \mathbb{R}^n$ (as a Lie group) and $Q = \mathbb{R}^n$ (as a manifold), with G acting on Q by translation. Eqs. (34) - (36) then yield the Poisson brackets $\{p_j, p_k\} = 0$, $\{p_j, q^k\} = \delta_j^k$, and $\{q^j, q^k\} = 0$, showing that in this case $M = \mathfrak{g}^* \times Q = \mathbb{R}^{2n}$ is the standard phase space of a particle moving in \mathbb{R}^n ; cf. (24). Second, the case $G = E(3)$ and $Q = \mathbb{R}^3$ yields a phase space $M = \mathbb{R}^3 \times \mathbb{R}^6$, where \mathbb{R}^6 is the phase space of a spinless particle just considered, and \mathbb{R}^3 is an additional internal space containing spin as a classical degree of freedom. Indeed, beyond the Poisson brackets on \mathbb{R}^6 just described, (34) - (36) give rise to the additional Poisson brackets $\{J_i, J_j\} = \epsilon_{ijk} J_k$, $\{J_i, p_j\} = \epsilon_{ijk} p_k$, and $\{J_i, q^j\} = \epsilon_{ijk} q^k$.¹⁵⁵

The analogy between (34), (35), (36) on the one hand, and (6), (8), (9), respectively, on the other, is no accident: the Poisson brackets in question *are* the

¹⁵⁵These are the classical counterparts of the commutation relations for angular momentum written down in footnote 108.

classical counterpart of the commutation relations just referred to. This observation is made precise by the fundamental theorem relating Mackey’s systems of imprimitivity to deformation quantization [Landsman, 1993; 1998]: one can equip the family of C^* -algebras

$$(37) \quad \begin{aligned} \mathcal{A}_0 &= C_0(\mathfrak{g}^* \times Q); \\ \mathcal{A}_\hbar &= C^*(G, Q), \end{aligned}$$

where $C^*(G, Q)$ is the transformation group C^* -algebra defined by the given G -action on Q (cf. the end of Subsection 4.1), with the structure of a continuous field, and one can define quantization maps $\mathcal{Q}_\hbar : C_c^\infty(\mathfrak{g}^* \times Q) \rightarrow C^*(G, Q)$ so as to obtain a deformation quantization of the phase space $\mathfrak{g}^* \times Q$. It turns out that for special functions of the type $X, Y \in \mathfrak{g}$, and $f = f(q)$ just considered, the equality

$$(38) \quad \frac{i}{\hbar} [\mathcal{Q}_\hbar(f), \mathcal{Q}_\hbar(g)] - \mathcal{Q}_\hbar(\{f, g\}) = 0$$

holds exactly (and not merely asymptotically for $\hbar \rightarrow 0$, as required in the fundamental axiom (25) for deformation quantization).

This result clarifies the status of Mackey’s quantization by systems of imprimitivity. The classical theory underlying the relations (4) is not the usual phase space T^*Q of a structureless particle moving on Q , but $M = \mathfrak{g}^* \times Q$. For simplicity we restrict ourselves to the transitive case $Q = G/H$ (with canonical left G -action). Then M coincides with T^*Q only when $H = \{e\}$ and hence $Q = G$; ¹⁵⁶ in general, the phase space $\mathfrak{g}^* \times (G/H)$ is *locally* of the form $T^*(G/H) \times \mathfrak{h}^*$ (where \mathfrak{h}^* is the dual of the Lie algebra of H). The internal degree of freedom described by \mathfrak{h}^* is a generalization of classical spin, which, as we have seen, emerges in the case $G = E(3)$ and $H = SO(3)$. All this is merely a special case of a vast class of deformation quantizations described by Lie groupoids; see [Landsman, 1998; 1999b; 2006a] and [Landsman and Ramazan, 2001]. ¹⁵⁷

4.4 Geometric quantization

Because of its use of abstract C^* -algebras, deformation quantization is a fairly sophisticated and recent technique. Historically, it was preceded by a more concrete and traditional approach called *geometric quantization*. ¹⁵⁸ Here the goal is to

¹⁵⁶For a Lie group G one has $T^*G \cong \mathfrak{g}^* \times G$.

¹⁵⁷A similar analysis can be applied to Isham’s [1984] quantization scheme mentioned in footnote 117. The unitary irreducible representations of the canonical group G_c stand in bijective correspondence with the nondegenerate representations of the group C^* -algebra $C^*(G_c)$ [Pedersen, 1979], which is a deformation quantization of the Poisson manifold \mathfrak{g}_c^* (i.e. the dual of the Lie algebra of G_c). This Poisson manifold contains the coadjoint orbits of G_c as “irreducible” classical phase spaces, of which only one is the cotangent bundle $T^*(G/H)$ one initially thought one was quantizing (see [Landsman, 1998] for the classification of the coadjoint orbits of semidirect products). All other orbits are mere lumber that one should avoid. See also [Robson, 1996]. If one is ready for groupoids, there is no need for the canonical group approach.

¹⁵⁸Geometric quantization was independently introduced by Kostant [1970] and Souriau [1969]. Major later treatments on the basis of the original formalism are [Guillemin and Sternberg, 1977;

firstly “quantize” a phase space M by a concretely given Hilbert space $\mathcal{H}(M)$, and secondly to map the classical observables (i.e. the real-valued smooth functions on M) into self-adjoint operators on \mathcal{H} (which after all play the role of observables in von Neumann’s formalism of quantum mechanics).¹⁵⁹ In principle, this program should align geometric quantization much better with the fundamental role unbounded self-adjoint operators play in quantum mechanics than deformation quantization, but in practice geometric quantization continues to be plagued by problems.¹⁶⁰ However, it would be wrong to see deformation quantization and geometric quantization as *competitors*; as we shall see in the next subsection, they are natural *allies*, forming “complementary” parts of a conjectural quantization functor.

In fact, in our opinion geometric quantization is best compared and contrasted with phase space quantization in its concrete formulation of Subsection 4.2 (i.e. before its C^* -algebraic abstraction and subsequent absorption into deformation quantization as indicated in Subsection 4.3).¹⁶¹ For geometric quantization equally well starts with the Hilbert space $L^2(M)$,¹⁶² and subsequently attempts to construct $\mathcal{H}(M)$ from it, though typically in a different way from (14).

Before doing so, however, the geometric quantization procedure first tries to define a linear map \mathcal{Q}_\hbar^{pre} from $C^\infty(M)$ to the class of (generally unbounded) operators on $L^2(M)$ that formally satisfies

$$(39) \quad \frac{i}{\hbar}[\mathcal{Q}_\hbar^{pre}(f), \mathcal{Q}_\hbar^{pre}(g)] - \mathcal{Q}_\hbar^{pre}(\{f, g\}) = 0,$$

i.e. (38) with $\mathcal{Q} = \mathcal{Q}_\hbar^{pre}$, as well as the nondegeneracy property

$$(40) \quad \mathcal{Q}_\hbar^{pre}(\chi_M) = 1,$$

where χ_M is the function on M that is identically equal to 1, and the 1 on the right-hand side is the unit operator on $L^2(M)$. Such a map is called a *prequantization*

Śniatycki, 1980; Kirillov, 1990; Woodhouse, 1992; Puta, 1993; Chernoff, 1995; Kirillov, 2004], and [Ali and Englis, 2004]. The modern era (based on the use of Dirac operators and K -theory) was initiated by unpublished remarks by Bott in the early 1990s; see [Vergne, 1994] and [Guillemin *et al.*, 2002]. The postmodern (i.e. functorial) epoch was launched in [Landsman, 2005].

¹⁵⁹In geometric quantization phase spaces are always seen as symplectic manifolds (with the sole exception of [Vaisman, 1991]; the reason why it is unnatural to start with the more general class of Poisson manifolds will become clear in the next subsection.

¹⁶⁰Apart from rather technical issues concerning the domains and self-adjointness properties of the operators defined by geometric quantization, the main point is that the various mathematical choices one has to make in the geometric quantization procedure cannot all be justified by physical arguments, although the physical properties of the theory depend on these choices. (The notion of a polarization is the principal case in point; see also footnote 173 below.) Furthermore, as we shall see, one cannot quantize sufficiently many functions in standard geometric quantization. Our functorial approach to geometric quantization in Subsection 4.5 was partly invented to alleviate these problems.

¹⁶¹See also [Tuynman, 1987].

¹⁶²Defined with respect to the Liouville measure times a suitable factor c_\hbar , as in (17) etc.; in geometric quantization this factor is not very important, as it is unusual to study the limit $\hbar \rightarrow 0$. For $M = \mathbb{R}^{2n}$ the measure on M with respect to which $L^2(M)$ is defined is $d^n p d^n q / (2\pi\hbar)^n$.

of M .¹⁶³ For $M = \mathbb{R}^{2n}$ (equipped with its standard Poisson bracket (24)), a prequantization map is given (on $\Phi \in L^2(M)$) by

$$(41) \quad \mathcal{Q}_\hbar^{pre}(f)\Phi = -i\hbar\{f, \Phi\} + \left(f - \sum_j p_j \frac{\partial f}{\partial p_j} \right) \Phi.$$

This expression is initially defined for $\Phi \in C_c^\infty(M) \subset L^2(M)$, on which domain $\mathcal{Q}_\hbar^{pre}(f)$ is symmetric when f is real-valued;¹⁶⁴ note that the operator in question is unbounded even when f is bounded.¹⁶⁵ This looks complicated; the simpler expression $\mathcal{Q}_\hbar(f)\Phi = -i\hbar\{f, \Phi\}$, however, would satisfy (38) but not (40), and the goal of the second term in (41) is to satisfy the latter condition while preserving the former.¹⁶⁶ For example, one has

$$(42) \quad \begin{aligned} \mathcal{Q}_\hbar^{pre}(q^k) &= q^k + i\hbar \frac{\partial}{\partial p_k}; \\ \mathcal{Q}_\hbar^{pre}(p_j) &= -i\hbar \frac{\partial}{\partial q^j}. \end{aligned}$$

For general phase spaces M one may construct a map \mathcal{Q}_\hbar^{pre} that satisfies (39) and (40) when M is “prequantizable”; a full explanation of this notion requires some differential geometry.¹⁶⁷ Assuming this to be the case, then for one thing prequantization is a very effective tool in constructing unitary group representations of the kind that are interesting for physics. Namely, suppose a Lie group G acts on the phase space M in “canonical” fashion. This means that there exists a map $\mu : M \rightarrow \mathfrak{g}^*$, called the *momentum map*, such that $\xi_{\mu_X} = \xi_X^M$ for each

¹⁶³The idea of prequantization predates geometric quantization; see [van Hove, 1951] and [Segal, 1960].

¹⁶⁴An operator A defined on a dense subspace $\mathcal{D} \subset \mathcal{H}$ of a Hilbert space \mathcal{H} is called *symmetric* when $(A\Psi, \Phi) = (\Psi, A\Phi)$ for all $\Psi, \Phi \in \mathcal{D}$.

¹⁶⁵As mentioned, self-adjointness is a problem in geometric quantization; we will not address this issue here. Berezin quantization is much better behaved than geometric quantization in this respect, since it maps bounded functions into bounded operators.

¹⁶⁶One may criticize the geometric quantization procedure for emphasizing (39) against its equally natural counterpart $\mathcal{Q}(fg) = \mathcal{Q}(f)\mathcal{Q}(g)$, which fails to be satisfied by \mathcal{Q}_\hbar^{pre} (and indeed by any known quantization procedure, except the silly $\mathcal{Q}(f) = f$ (as a multiplication operator on $L^2(M)$)).

¹⁶⁷A symplectic manifold (M, ω) is called *prequantizable* at some fixed value of \hbar when it admits a complex line bundle $L \rightarrow M$ (called the *prequantization line bundle*) with connection ∇ such that $F = -i\omega/\hbar$ (where F is the curvature of the connection, defined by $F(X, Y) = [\nabla_X, \nabla_Y] - \nabla_{[X, Y]}$). This is the case iff $[\omega]/2\pi\hbar \in H^2(M, \mathbb{Z})$, where $[\omega]$ is the de Rham cohomology class of the symplectic form. If so, prequantization is defined by the formula $\mathcal{Q}_\hbar^{pre}(f) = -i\hbar\nabla_{\xi_f} + f$, where ξ_f is the Hamiltonian vector field of f (see Subsection 4.3). This expression is defined and symmetric on the space $C_c^\infty(M, L) \subset L^2(M)$ of compactly supported smooth sections of L , and is easily checked to satisfy (39) and (40). To obtain (41) as a special case, note that for $M = \mathbb{R}^{2n}$ with the canonical symplectic form $\omega = \sum_k dp_k \wedge dq^k$ one has $[\omega] = 0$, so that L is the trivial bundle $L = \mathbb{R}^{2n} \times \mathbb{C}$. The connection $\nabla = d + A$ with $A = -\frac{i}{\hbar} \sum_k p_k dq^k$ satisfies $F = -i\omega/\hbar$, and this eventually yields (41).

$X \in \mathfrak{g}$,¹⁶⁸ and in addition $\{\mu_X, \mu_Y\} = \mu_{[X, Y]}$. See [Abraham and Marsden, 1985; Marsden and Ratiu, 1994; Landsman, 1998; Belot, 2005; Butterfield, 2005], etc. On then obtains a representation π of the Lie algebra \mathfrak{g} of G by skew-symmetric unbounded operators on $L^2(M)$ through

$$(43) \quad \pi(X) = -i\hbar \mathcal{Q}_\hbar^{pre}(\mu_X),$$

which often exponentiates to a unitary representation of G .¹⁶⁹

As the name suggests, prequantization is not yet quantization. For example, the prequantization of $M = \mathbb{R}^{2n}$ does not reproduce Schrödinger’s wave mechanics: the operators (42) are not unitarily equivalent to (2). In fact, as a carrier of the representation (42) of the canonical commutation relations (1), the Hilbert space $L^2(\mathbb{R}^{2n})$ contains $L^2(\mathbb{R}^n)$ (carrying the representation (2)) with infinite multiplicity [Ali and Emch, 1986]. This situation is often expressed by the statement that “prequantization is reducible” or that the prequantization Hilbert space $L^2(M)$ is ‘too large’, but both claims are misleading: $L^2(M)$ is actually *irreducible* under the action of $\mathcal{Q}_\hbar^{pre}(C^\infty(M))$ [Tuynman, 1998], and saying that for example $L^2(\mathbb{R}^n)$ is “larger” than $L^2(\mathbb{R}^n)$ is unmathematical in view of the unitary isomorphism of these Hilbert spaces. What is really meant here is that in typical examples $L^2(M)$ is generically reducible under the action of some Lie algebra where one would like it to be irreducible. This applies, for example, to (2), which defines a representation of the Lie algebra of the Heisenberg group. More generally, in the case where a phase space M carries a transitive action of a Lie group G , so that one would expect the quantization of this G -action by unitary operators on a Hilbert space to be irreducible, $L^2(M)$ is typically highly reducible under the representation (43) of \mathfrak{g} .¹⁷⁰

Phase space quantization encounters this problem as well. Instead of the complicated expression (41), through (11) it simply “phase space prequantizes” $f \in C^\infty(M)$ on $L^2(M)$ by f as a multiplication operator.¹⁷¹ Under this action of $C^\infty(M)$ the Hilbert space $L^2(M)$ is of course highly reducible.¹⁷² The identifica-

¹⁶⁸Here $\mu_X \in C^\infty(M)$ is defined by $\mu_X(x) = \langle \mu(x), X \rangle$, and ξ_X^M is the vector field on M defined by the G -action (cf. footnote 107). Hence this condition means that $\{\mu_X, f\}(y) = d/dt|_{t=0}[f(\exp(-tX)y)]$ for all $f \in C^\infty(M)$ and all $y \in M$.

¹⁶⁹An operator A defined on a dense subspace $\mathcal{D} \subset \mathcal{H}$ of a Hilbert space \mathcal{H} is called *skew-symmetric* when $(A\Psi, \Phi) = -(\Psi, A\Phi)$ for all $\Psi, \Phi \in \mathcal{D}$. If one has a unitary representation U of a Lie group G on \mathcal{H} , then the derived representation dU of the Lie algebra \mathfrak{g} (see footnote 104) consists of skew-symmetric operators, making one hopeful that a given representation of \mathfrak{g} by skew-symmetric operators can be integrated (or exponentiated) to a unitary representation of G . See [Barut and Raçka, 1977] or [Jørgensen and Moore, 1984] and references therein.

¹⁷⁰This can be made precise in the context of the so-called orbit method, cf. the books cited in footnote 158.

¹⁷¹For unbounded f this operator is defined on the set of all $\Phi \in L^2(M)$ for which $f\Phi \in L^2(M)$.

¹⁷²Namely, each (measurable) subset $E \subset M$ defines a projection χ_E , and $\chi_E L^2(M)$ is stable under all multiplication operators f . One could actually decide not to be bothered by this problem and stop here, but then one is simply doing classical mechanics in a Hilbert space setting [Koopman, 1931]. This formalism even turns out to be quite useful for ergodic theory [Reed and Simon, 1972].

tion of an appropriate subspace

$$(44) \quad \mathcal{H}(M) = pL^2(M)$$

of $L^2(M)$ (where p is a projection) as the Hilbert space carrying the “quantization” of M (or rather of $C^\infty(M)$) may be seen as a solution to this reducibility problem, for if the procedure is successful, the projection p is chosen such that $pL^2(M)$ is irreducible under $pC^\infty(M)p$. Moreover, in this way practically any function on M can be quantized, albeit at the expense of (38) (which, as we have seen, gets replaced by its asymptotic version (25)). See Subsection 6.3 for a discussion of reducibility versus irreducibility of representations of algebras of observables in classical and quantum theory.

We restrict our treatment of geometric quantization to situations where it adopts the same strategy as above, in assuming that the final Hilbert space has the form (44) as well.¹⁷³ But it crucially differs from phase space quantization in that its first step is (41) (or its generalization to more general phase spaces) rather than just having $f\Phi$ on the right-hand side.¹⁷⁴ Moreover, in geometric quantization one merely quantizes a *subspace* of the set $C^\infty(M)$ of classical observables, consisting of those functions that satisfy

$$(45) \quad [\mathcal{Q}_\hbar^{pre}(f), p] = 0.$$

If a function $f \in C^\infty(M)$ satisfies this condition, then one defines the “geometric quantization” of f as

$$(46) \quad \mathcal{Q}_\hbar^G(f) = \mathcal{Q}_\hbar^{pre}(f) \upharpoonright \mathcal{H}(M).$$

This is well defined, since because of (45) the operator $\mathcal{Q}_\hbar^{pre}(f)$ now maps $pL^2(M)$ onto itself. Hence (38) holds for $\mathcal{Q}_\hbar = \mathcal{Q}_\hbar^G$ because of (39); in geometric quantization one simply refuses to quantize functions for which (38) is *not* valid.

Despite some impressive initial triumphs,¹⁷⁵ there is no general method that accomplishes the goals of geometric quantization with guaranteed success. Therefore, geometric quantization has remained something like a hacker’s tool, whose applicability largely depends on the creativity of the user.

In any case, our familiar example $M = \mathbb{R}^{2n}$ is well understood, and we illustrate the general spirit of the method in its setting, simplifying further by taking $n = 1$. It is convenient to replace the canonical coordinates (p, q) on M by $z = p + iq$ and $\bar{z} = p - iq$, and the mathematical toolkit of geometric quantization makes it very

¹⁷³ Geometric quantization has traditionally been based on the notion of a polarization (cf. the references in footnote 158). This device produces a final Hilbert space $\mathcal{H}(M)$ which may not be a subspace of $L^2(M)$, except in the so-called (anti-) holomorphic case.

¹⁷⁴It also differs from phase space quantization in the ideology that the projection p ought to be constructed solely from the geometry of M : hence the name ‘geometric quantization’.

¹⁷⁵Such as the orbit method for nilpotent groups and the newly understood Borel–Weil method for compact groups, cf. [Kirillov, 2004] and most other books cited in footnote 158.

natural to look at the space of solutions within $L^2(\mathbb{R}^2)$ of the equations¹⁷⁶

$$(47) \quad \left(\frac{\partial}{\partial \bar{z}} + \frac{z}{4\hbar} \right) \Phi(z, \bar{z}) = 0.$$

The general solution of these equations that lies in $L^2(\mathbb{R}^2) = L^2(\mathbb{C})$ is

$$(48) \quad \Phi(z, \bar{z}) = e^{-|z|^2/4\hbar} f(z),$$

where f is a holomorphic function such that

$$(49) \quad \int_{\mathbb{C}} \frac{dz d\bar{z}}{2\pi\hbar i} e^{-|z|^2/2\hbar} |f(z)|^2 < \infty.$$

The projection p , then, is the projection onto the closed subspace of $L^2(\mathbb{C})$ consisting of these solutions.¹⁷⁷ The Hilbert space $pL^2(\mathbb{C})$ is unitarily equivalent to $L^2(\mathbb{R})$ in a natural way (i.e. without the choice of a basis). The condition (45) boils down to $\partial^2 f(z, \bar{z})/\partial \bar{z}_i \partial \bar{z}_j = 0$; in particular, the coordinate functions q and p are quantizable. Transforming to $L^2(\mathbb{R})$, one finds that the operators $\mathcal{Q}_\hbar^G(q)$ and $\mathcal{Q}_\hbar^G(p)$ coincide with Schrödinger’s expressions (2). In particular, the Heisenberg group H_1 , which acts with infinite multiplicity on $L^2(\mathbb{C})$, acts irreducibly on $pL^2(\mathbb{C})$.

4.5 Epilogue: functoriality of quantization

A very important aspect of quantization is its interplay with symmetries and constraints. Indeed, the fundamental theories describing Nature (viz. electrodynamics, Yang–Mills theory, general relativity, and possibly also string theory) are a priori formulated as constrained systems. The classical side of constraints and reduction is well understood,¹⁷⁸ a large class of important examples being codified by the procedure of symplectic reduction. A special case of this is *Marsden–Weinstein reduction*: if a Lie group G acts on a phase space M in canonical fashion with momentum map $\mu : M \rightarrow \mathfrak{g}^*$ (cf. Subsection 4.4), one may form another phase space $M//G = \mu^{-1}(0)/G$.¹⁷⁹ Physically, in the case where G is a gauge group and M is the unconstrained phase space, $\mu^{-1}(0)$ is the constraint hypersurface (i.e. the

¹⁷⁶Using the formalism explained in footnote 167, we replace the 1-form $A = -\frac{i}{\hbar} \sum_k p_k dq^k$ defining the connection $\nabla = d + A$ by the gauge-equivalent form $A = \frac{i}{2\hbar} (\sum_k q^k dp_k - \sum_k p_k dq^k) = -\frac{i}{\hbar} \sum_k p_k dq^k + \frac{i}{2\hbar} d(\sum_k p_k q^k)$, which has the same curvature. In terms of this new A , which in complex coordinates reads $A = \sum_k (z_k d\bar{z}_k - \bar{z}_k dz_k)/4\hbar$, eq. (47) is just $\nabla_{\partial/\partial \bar{z}} \Phi = 0$. This is an example of the so-called holomorphic polarization in the formalism of geometric quantization.

¹⁷⁷The collection of all holomorphic functions on \mathbb{C} satisfying (49) is a Hilbert space with respect to the inner product $(f, g) = (2\pi\hbar i)^{-1} \int_{\mathbb{C}} dz d\bar{z} \exp(-|z|^2/2\hbar) \bar{f}(\bar{z})g(z)$, called the *Bargmann–Fock space* \mathcal{H}_{BF} . This space may be embedded in $L^2(\mathbb{C})$ by $f(z) \mapsto \exp(-|z|^2/2\hbar)f(z)$, and the image of this embedding is of course just $pL^2(\mathbb{C})$.

¹⁷⁸See [Gotay *et al.*, 1978; Binz *et al.*, 1988; Marsden, 1992; Marsden and Ratiu, 1994; Landsman, 1998; Butterfield, 2005], and [Belot, 2005].

¹⁷⁹Technically, M has to be a symplectic manifold, and if G acts properly and freely on $\mu^{-1}(0)$, then $M//G$ is again a symplectic manifold.

subspace of M on which the constraints defined by the gauge symmetry hold), and $M//G$ is the true phase space of the system that only contains physical degrees of freedom.

Unfortunately, the correct way of dealing with constrained quantum systems remains a source of speculation and controversy.¹⁸⁰ practically all rigorous results on quantization (like the ones discussed in the preceding subsections) concern unconstrained systems. Accordingly, one would like to quantize a constrained system by reducing the problem to the unconstrained case. This could be done provided the following scenario applies. One first quantizes the unconstrained phase space M (supposedly the easiest part of the problem), and subsequently imposes a quantum version of symplectic reduction. Finally, one proves by abstract means that the quantum theory thus constructed is equal to the theory defined by first reducing at the classical level and then quantizing the constrained classical phase space (usually an impossible task to perform in practice).

Tragically, sufficiently powerful theorems stating that “quantization commutes with reduction” in this sense remain elusive.¹⁸¹ So far, this has blocked, for example, a rigorous quantization of Yang–Mills theory in dimension 4; this is one of the Millennium Problems of the Clay Mathematical Institute, rewarded with 1 Million dollars.¹⁸²

On a more spiritual note, the mathematician E. Nelson famously said that ‘First quantization is a mystery, but second quantization is a functor.’ The functoriality of ‘second’ quantization (a construction involving Fock spaces, see [Reed and Simon, 1975] being an almost trivial matter, the deep mathematical and conceptual problem lies in the possible functoriality of ‘first’ quantization, which simply means quantization in the sense we have been discussing so far. This was initially taken to mean that canonical transformations α of the phase space M should be ‘quantized’ by unitary operators $U(\alpha)$ on $\mathcal{H}(M)$, in such a way $U(\alpha)\mathcal{Q}_\hbar(f)U(\alpha)^{-1} = \mathcal{Q}(L_\alpha f)$ (cf. (10)). This is possible only in special circumstances, e.g., when $M = \mathbb{R}^{2n}$ and α is a linear symplectic map, and more generally when $M = G/H$ is homogeneous and $\alpha \in G$ (see the end of Subsection 4.2).¹⁸³ Consequently, the functoriality of quantization is widely taken to be a dead end.¹⁸⁴

However, all no-go theorems establishing this conclusion start from wrong and

¹⁸⁰Cf. [Dirac, 1964; Sundermeyer, 1982; Gotay, 1986; Duval *et al.*, 1991; Govaerts, 1991; Henneaux and Teitelboim, 1992], and [Landsman, 1998] for various perspectives on the quantization of constrained systems.

¹⁸¹The so-called Guillemin–Sternberg conjecture [Guillemin and Sternberg, 1982] — now a theorem [Meinrenken, 1998; Meinrenken and Sjamaar, 1999] — merely deals with the case of Marsden–Weinstein reduction where G and M are compact. Mathematically impressive as the “quantization commutes with reduction” theorem already is here, it is a far call from gauge theories, where the groups and spaces are not only noncompact but even infinite-dimensional.

¹⁸²See <http://www.claymath.org/millennium/>

¹⁸³Canonical transformations *can* be quantized in approximate sense that becomes precise as $\hbar \rightarrow 0$ by means of so-called Fourier integral operators; see [Hörmander, 1971; 1985b] and [Duistermaat, 1996].

¹⁸⁴See [Groenewold, 1946; van Hove, 1951; Gotay *et al.*, 1996], and [Gotay, 1999].

naive categories, both on the classical and on the quantum side.¹⁸⁵ It appears very likely that one may indeed make quantization functorial by a more sophisticated choice of categories, with the additional bonus that deformation quantization and geometric quantization become unified: the former is the object part of the quantization functor, whereas the latter (suitably reinterpreted) is the arrow part. Amazingly, on this formulation the statement that ‘quantization commutes with reduction’ becomes a special case of the functoriality of quantization [Landsman, 2002; 2005].

To explain the main idea, we return to the geometric quantization of $M = \mathbb{R}^2 \cong \mathbb{C}$ explained in the preceding subsection. The identification of $pL^2(\mathbb{C})$ ¹⁸⁶ as the correct Hilbert space of the problem may be understood in a completely different way, which paves the way for the powerful reformulation of the geometric quantization program that will eventually define the quantization functor. Namely, \mathbb{C} supports a certain linear first-order differential operator \mathcal{D} that is entirely defined by its geometry as a phase space, called the *Dirac operator*.¹⁸⁷ This operator is given by¹⁸⁸

$$(50) \quad \mathcal{D} = 2 \begin{pmatrix} 0 & -\frac{\partial}{\partial z} + \frac{\bar{z}}{4\hbar} \\ \frac{\partial}{\partial \bar{z}} + \frac{z}{4\hbar} & 0 \end{pmatrix},$$

acting on $L^2(\mathbb{C}) \otimes \mathbb{C}^2$ (as a suitably defined unbounded operator). This operator has the generic form

$$\mathcal{D} = \begin{pmatrix} 0 & \mathcal{D}_- \\ \mathcal{D}_+ & 0 \end{pmatrix}.$$

The *index* of such an operator is given by

$$(51) \quad \text{index}(\mathcal{D}) = [\ker(\mathcal{D}_+)] - [\ker(\mathcal{D}_-)],$$

where $[\ker(\mathcal{D}_\pm)]$ stand for the (unitary) isomorphism class of $\ker(\mathcal{D}_\pm)$ *seen as a representation space of a suitable algebra of operators*.¹⁸⁹ In the case at hand,

¹⁸⁵Typically, one takes the classical category to consist of symplectic manifolds as objects and symplectomorphisms as arrows, and the quantum category to have C^* -algebras as objects and automorphisms as arrows.

¹⁸⁶Or the Bargmann–Fock space \mathcal{H}_{BF} , see footnote 177.

¹⁸⁷Specifically, this is the so-called Spin^c Dirac operator defined by the complex structure of \mathbb{C} , *coupled to the prequantization line bundle*. See [Guillemin *et al.*, 2002].

¹⁸⁸Relative to the Dirac matrices $\gamma^1 = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$ and $\gamma^2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$.

¹⁸⁹The left-hand side of (51) should really be written as $\text{index}(\mathcal{D}_+)$, since $\text{coker}(\mathcal{D}_+) = \ker(\mathcal{D}_+^*)$ and $\mathcal{D}_+^* = \mathcal{D}_-$, but since the index is naturally associated to \mathcal{D} as a whole, we abuse notation in writing $\text{index}(\mathcal{D})$ for $\text{index}(\mathcal{D}_+)$. The usual index of a linear map $L : V \rightarrow W$ between finite-dimensional vector spaces is defined as $\text{index}(L) = \dim(\ker(L)) - \dim(\text{coker}(L))$, where $\text{coker}(L) = W/\text{ran}(L)$. Elementary linear algebra yields $\text{index}(L) = \dim(V) - \dim(W)$. This is surprising because it is independent of L , whereas $\dim(\ker(L))$ and $\dim(\text{coker}(L))$ quite sensitively depend on it. For, example, take $V = W$ and $L = \varepsilon \cdot 1$. If $\varepsilon \neq 0$ then $\dim(\ker(\varepsilon \cdot 1)) = \dim(\text{coker}(\varepsilon \cdot 1)) = 0$, whereas for $\varepsilon = 0$ one has $\dim(\ker(0)) = \dim(\text{coker}(0)) = \dim(V)$! Similarly, the usual definition of geometric quantization through (47) etc. is unstable against perturbations of the underlying symplectic structure, whereas the refined definition through (51) is not. To

one has $\ker(\mathcal{D}_+) = pL^2(\mathbb{C})$ (cf. (47) etc.) and $\ker(\mathcal{D}_-) = 0$,¹⁹⁰ where we regard $\ker(\mathcal{D}_+)$ as a representation space of the Heisenberg group H_1 . Consequently, the geometric quantization of the phase space \mathbb{C} is given *modulo unitary equivalence* by $\text{index}(\mathcal{D})$, seen as a “formal difference” of representation spaces of H_1 .

This procedure may be generalized to arbitrary phase spaces M , where \mathcal{D} is a certain operator naturally defined by the phase space geometry of M and the demands of quantization.¹⁹¹ This has turned out to be the most promising formulation of geometric quantization — at some cost.¹⁹² For the original goal of quantizing a phase space by a Hilbert space has now been replaced by a much more abstract procedure, in which the result of quantization is a formal difference of certain isomorphism classes of representation spaces of the quantum algebra of observables. To illustrate the degree of abstraction involved here, suppose we ignore the action of the observables (such as position and momentum in the example just considered). In that case the isomorphism class $[\mathcal{H}]$ of a Hilbert space \mathcal{H} is entirely characterized by its dimension $\dim(\mathcal{H})$, so that (in case that $\ker(\mathcal{D}_-) \neq 0$) quantization (in the guise of $\text{index}(\mathcal{D})$) can even be a negative number! Have we gone mad?

Not quite. The above picture of geometric quantization is indeed quite irrelevant to physics, *unless it is supplemented by deformation quantization*. It is convenient to work at some fixed value of \hbar in this context, so that deformation quantization merely associates some C^* -algebra $\mathcal{A}(P)$ to a given phase space P .¹⁹³ Looking for a categorical interpretation of quantization, it is therefore natural to assume that the objects of the classical category \mathfrak{C} are phase spaces P ,¹⁹⁴ whereas the

pass to the latter from the above notion of an index, we first write $\text{index}(L) = [\ker(L)] - [\text{coker}(L)]$, where $[X]$ is the isomorphism class of a linear space X as a \mathbb{C} -module. This expression is an element of $K_0(\mathbb{C})$, and we recover the earlier index through the realization that the class $[X]$ is entirely determined by $\dim(X)$, along with and the corresponding isomorphism $K_0(\mathbb{C}) \cong \mathbb{Z}$. When a more complicated finite-dimensional C^* -algebra \mathcal{A} acts on V and W with the property that $\ker(L)$ and $\text{coker}(L)$ are stable under the \mathcal{A} -action, one may define $[\ker(L)] - [\text{coker}(L)]$ and hence $\text{index}(L)$ as an element of the so-called C^* -algebraic K-theory group $K_0(\mathcal{A})$. Under certain technical conditions, this notion of an index may be generalized to infinite-dimensional Hilbert spaces and C^* -algebras; see [Baum *et al.*, 1994] and [Blackadar, 1998]. The K-theoretic index is best understood when $\mathcal{A} = C^*(G)$ is the group C^* -algebra of some locally compact group G . In the example $M = \mathbb{R}^2$ one might take G to be the Heisenberg group H_1 , so that $\text{index}(\mathcal{D}) \in K_0(C^*(H_1))$. See [Elliott *et al.*, 1993] for a description of this K_0 -group.

¹⁹⁰Since $(-\frac{\partial}{\partial z} + \frac{\bar{z}}{4\hbar})\Phi = 0$ implies $\Phi(z, \bar{z}) = \exp(|z|^2/4\hbar)f(\bar{z})$, which lies in $L^2(\mathbb{C})$ iff $f = 0$.

¹⁹¹Any symplectic manifold carries an almost complex structure compatible with the symplectic form, leading to a Spin^c Dirac operator as described in footnote 187. See, again, [Guillemin *et al.*, 2002]. If $M = G/H$, or, more generally, if M carries a canonical action of a Lie group G with compact quotient M/G , then $\text{index}(\mathcal{D})$ defines an element of $K_0(C^*(G))$. See footnote 189. In complete generality, $\text{index}(\mathcal{D})$ ought to be an element of $K_0(\mathcal{A})$, where \mathcal{A} is the C^* -algebra of observables of the quantum system.

¹⁹²On the benefit side, the invariance of the index under continuous deformations of \mathcal{D} seems to obviate the ambiguity of traditional quantization procedures with respect to different ‘operator orderings’ not prescribed by the classical theory.

¹⁹³Here P is not necessarily symplectic; it may be a Poisson manifold, and to keep Poisson and symplectic manifolds apart we denote the former by P from now on, preserving the notation M for the latter.

¹⁹⁴Strictly speaking, to be an object in this category a Poisson manifold P must be *integrable*;

objects of the quantum category \mathfrak{Q} are C^* -algebras.¹⁹⁵ The object part of the hypothetical quantization functor is to be deformation quantization, symbolically written as $P \mapsto \mathcal{Q}(P)$.

Everything then fits together if geometric quantization is reinterpreted as the arrow part of the conjectural quantization functor. To accomplish this, the arrows in the classical category \mathfrak{C} should not be taken to be maps between phase spaces, but *symplectic bimodules* $P_1 \leftarrow M \rightarrow P_2$.¹⁹⁶ More precisely, the arrows in \mathfrak{C} are suitable isomorphism classes of such bimodules.¹⁹⁷ Similarly, the arrows in the quantum category \mathfrak{Q} are not morphisms of C^* -algebras, as might naively be expected, but certain isomorphism classes of bimodules for C^* -algebras, equipped with the additional structure of a generalized Dirac operator.¹⁹⁸

Having already defined the object part of the quantization map $\mathcal{Q} : \mathfrak{C} \rightarrow \mathfrak{Q}$ as deformation quantization, we now propose that the arrow part is geometric quantization, in the sense of a suitable generalization of (51); see [Landsman, 2005] for details. We then conjecture that \mathcal{Q} is a functor; in the cases where this can and has been checked, the functoriality of \mathcal{Q} is precisely the statement that quantization commutes with reduction.¹⁹⁹

Thus Heisenberg's idea of *Umdeutung* finds its ultimate realization in the quantization functor.

5 THE LIMIT $\hbar \rightarrow 0$

It was recognized at an early stage that the limit $\hbar \rightarrow 0$ of Planck's constant going to zero should play a role in the explanation of the classical world from quantum theory. Strictly speaking, \hbar is a dimensionful *constant*, but in practice one studies the semiclassical regime of a given quantum theory by forming a dimensionless combination of \hbar and other parameters; this combination then re-enters the theory as if it were a dimensionless version of \hbar that can indeed be varied. The oldest example is Planck's radiation formula (1), with temperature T as the pertinent

see [Landsman, 2001].

¹⁹⁵For technical reasons involving K -theory these have to be separable.

¹⁹⁶Here M is a symplectic manifold and P_1 and P_2 are integrable Poisson manifolds; the map $M \rightarrow P_2$ is anti-Poisson, whereas the map $P_1 \leftarrow M$ is Poisson. Such bimodules (often called *dual pairs*) were introduced by Karasev [1989] and Weinstein [1983]. In order to occur as arrows in \mathfrak{C} , symplectic bimodules have to satisfy a number of regularity conditions [Landsman, 2001].

¹⁹⁷This is necessary in order to make arrow composition associative; this is given by a generalization of the symplectic reduction procedure.

¹⁹⁸The category \mathfrak{Q} is nothing but the category KK introduced by Kasparov, whose objects are separable C^* -algebras, and whose arrows are the so-called Kasparov group $KK(A, B)$, composed with Kasparov's product $KK(A, B) \times KK(B, C) \rightarrow KK(A, C)$. See [Higson, 1990] and [Blackadar, 1998].

¹⁹⁹A canonical G -action on a symplectic manifold M with momentum map $\mu : M \rightarrow \mathfrak{g}^*$ gives rise to a dual pair $pt \leftarrow M \rightarrow \mathfrak{g}^*$, which in \mathfrak{C} is interpreted as an arrow from the space pt with one point to \mathfrak{g}^* . The composition of this arrow with the arrow $\mathfrak{g}^* \leftarrow 0 \rightarrow pt$ from \mathfrak{g}^* to pt is $pt \leftarrow M//G \rightarrow pt$. If G is connected, functoriality of quantization on these two pairs is equivalent to the Guillemin–Sternberg conjecture (cf. footnote 181); see [Landsman, 2005].

variable. Indeed, the observation of Einstein [1905] and Planck [1906] that in the limit $\hbar\nu/kT \rightarrow 0$ this formula converges to the classical equipartition law $E_\nu/N_\nu = kT$ may well be the first use of the $\hbar \rightarrow 0$ limit of quantum theory.²⁰⁰

Another example is the Schrödinger equation (3) with Hamiltonian $H = -\frac{\hbar^2}{2m} \Delta_x + V(x)$, where m is the mass of the pertinent particle. Here one may pass to dimensionless parameters by introducing an energy scale ϵ typical of H , like $\epsilon = \sup_x |V(x)|$, as well as a typical length scale λ , such as $\lambda = \epsilon / \sup_x |\nabla V(x)|$ (if these quantities are finite). In terms of the dimensionless variable $\tilde{x} = x/\lambda$, the rescaled Hamiltonian $\tilde{H} = H/\epsilon$ is then dimensionless and equal to $\tilde{H} = -\tilde{\hbar}^2 \Delta_{\tilde{x}} + \tilde{V}(\tilde{x})$, where $\tilde{\hbar} = \hbar/\lambda\sqrt{2m\epsilon}$ and $\tilde{V}(\tilde{x}) = V(\lambda\tilde{x})/\epsilon$. Here $\tilde{\hbar}$ is dimensionless, and one might study the regime where it is small [Gustafson and Sigal, 2003]. Our last example will occur in the theory of large quantum systems, treated in the next Section. In what follows, whenever it is considered variable \hbar will denote such a dimensionless version of Planck's constant.

Although, as we will argue, the limit $\hbar \rightarrow 0$ cannot by itself explain the classical world, it does give rise to a number of truly pleasing mathematical results. These, in turn, render almost inescapable the conclusion that the limit in question is indeed a relevant one for the recovery of classical physics from quantum theory. Thus the present section is meant to be a catalogue of those pleasantries that might be of direct interest to researchers in the foundations of quantum theory.

There is another, more technical use of the $\hbar \rightarrow 0$ limit, which is to perform computations in quantum mechanics by approximating the time-evolution of states and observables in terms of associated classical objects. This endeavour is known as *semiclassical analysis*. Mathematically, this use of the $\hbar \rightarrow 0$ limit is closely related to the goal of recovering classical mechanics from quantum mechanics, but conceptually the matter is quite different. We will attempt to bring the pertinent differences out in what follows.

5.1 Coherent states revisited

As Schrödinger [1926b] foresaw, coherent states play an important role in the limit $\hbar \rightarrow 0$. We recall from Subsection 4.2 that *for some fixed value \hbar of Planck's constant* coherent states in a Hilbert space \mathcal{H} for a phase space M are defined by an injection $M \hookrightarrow \mathcal{H}$, $z \mapsto \Psi_z^\hbar$, such that (16) and (17) hold. In what follows, we shall say that Ψ_z^\hbar is *centered at $z \in M$* , a terminology justified by the key example (20).

To be relevant to the classical limit, coherent states must satisfy an additional property concerning their dependence on \hbar , which also largely clarifies their nature [Landsman, 1998]. Namely, we require that for each $f \in C_c(M)$ and each $z \in M$ the following function from the set I in which \hbar takes values (i.e. usually $I = [0, 1]$,

²⁰⁰Here Einstein [1905] put $\hbar\nu/kT \rightarrow 0$ by letting $\nu \rightarrow 0$ at fixed T and \hbar , whereas Planck [1906] took $T \rightarrow \infty$ at fixed ν and \hbar .

but in any case containing zero as an accumulation point) to \mathbb{C} is continuous:

- (1) $\hbar \mapsto c_{\hbar} \int_M d\mu_L(w) |(\Psi_w^{\hbar}, \Psi_z^{\hbar})|^2 f(w) \quad (\hbar > 0);$
- (2) $0 \mapsto f(z).$

In view of (19), the right-hand side of (2) is the same as $(\Psi_z^{\hbar}, \mathcal{Q}_{\hbar}^B(f)\Psi_z^{\hbar})$. In particular, this continuity condition implies

$$(3) \quad \lim_{\hbar \rightarrow 0} (\Psi_z^{\hbar}, \mathcal{Q}_{\hbar}^B(f)\Psi_z^{\hbar}) = f(z).$$

This means that the classical limit of the quantum-mechanical expectation value of the phase space quantization (19) of the classical observable f in a coherent state centered at $z \in M$ is precisely the classical expectation value of f in the state z . This interpretation rests on the identification of classical states with probability measures on phase space M , under which points of M in the guise of Dirac measures (i.e. delta functions) are pure states. Furthermore, it can be shown (cf. [Landsman, 1998]) that the continuity of all functions (1) - (2) implies the property

$$(4) \quad \lim_{\hbar \rightarrow 0} |(\Psi_w^{\hbar}, \Psi_z^{\hbar})|^2 = \delta_{wz},$$

where δ_{wz} is the ordinary Kronecker delta (i.e. $\delta_{wz} = 0$ whenever $w \neq z$ and $\delta_{zz} = 1$ for all $z \in M$). This has a natural physical interpretation as well: the classical limit of the quantum-mechanical transition probability between two coherent states centered at $w, z \in M$ is equal to the classical (and trivial) transition probability between w and z . In other words, when \hbar becomes small, coherent states at different values of w and z become increasingly orthogonal to each other.²⁰¹ This has the interesting consequence that

$$(5) \quad \lim_{\hbar \rightarrow 0} (\Psi_w^{\hbar}, \mathcal{Q}_{\hbar}^B(f)\Psi_z^{\hbar}) = 0 \quad (w \neq z).$$

for all $f \in C_c(M)$. In particular, the following phenomenon of the Schrödinger cat type occurs in the classical limit: if $w \neq z$ and one has continuous functions $\hbar \mapsto c_w^{\hbar} \in \mathbb{C}$ and $\hbar \mapsto c_z^{\hbar} \in \mathbb{C}$ on $\hbar \in [0, 1]$ such that

$$(6) \quad \Psi_{w,z}^{\hbar} = c_w^{\hbar} \Psi_w^{\hbar} + c_z^{\hbar} \Psi_z^{\hbar}$$

is a unit vector for $\hbar \geq 0$ and also $|c_w^0|^2 + |c_z^0|^2 = 1$, then

$$(7) \quad \lim_{\hbar \rightarrow 0} (\Psi_{w,z}^{\hbar}, \mathcal{Q}_{\hbar}^B(f)\Psi_{w,z}^{\hbar}) = |c_w^0|^2 f(w) + |c_z^0|^2 f(z).$$

Hence the family of (typically) pure states $\psi_{w,z}^{\hbar}$ (on the C^* -algebras \mathcal{A}_{\hbar} in which the map \mathcal{Q}_{\hbar}^B takes values)²⁰² defined by the vectors $\Psi_{w,z}^{\hbar}$ in some sense converges

²⁰¹See [Mielnik, 1968; Cantoni, 1975; Beltrametti and Cassinelli, 1984; Landsman, 1998], and Subsection 6.3 below for the general meaning of the concept of a transition probability.

²⁰²For example, for $M = \mathbb{R}^{2n}$ each \mathcal{A}_{\hbar} is equal to the C^* -algebra of compact operators on $L^2(\mathbb{R}^n)$, on which each vector state is certainly pure.

to the mixed state on $C_0(M)$ defined by the right-hand side of (7). This is made precise at the end of this subsection.

It goes without saying that Schrödinger’s coherent states (20) satisfy our axioms; one may also verify (4) immediately from (21). Consequently, by (32) one has the same property (3) for Weyl quantization (as long as $f \in \mathcal{S}(\mathbb{R}^{2n})$),²⁰³ that is,

$$(8) \quad \lim_{\hbar \rightarrow 0} (\Psi_z^\hbar, \mathcal{Q}_\hbar^W(f), \Psi_z^\hbar) = f(z).$$

Similarly, (5) holds for \mathcal{Q}_\hbar^W as well.

In addition, many constructions referred to as coherent states in the literature (cf. the references in footnote 129) satisfy (16), (17), and (4); see [Landsman, 1998].²⁰⁴ The general picture that emerges is that a coherent state centered at $z \in M$ is the *Umdeutung* of z (seen as a classical pure state, as explained above) as a quantum-mechanical pure state.²⁰⁵

Despite their wide applicability (and some would say beauty), one has to look beyond coherent states for a complete picture of the $\hbar \rightarrow 0$ limit of quantum mechanics. The appropriate generalization is the concept of a *continuous field of states*.²⁰⁶ This is defined relative to a given deformation quantization of a phase space M ; cf. Subsection 4.3. If one now has a state ω_\hbar on \mathcal{A}_\hbar for each $\hbar \in [0, 1]$ (or, more generally, for a discrete subset of $[0, 1]$ containing 0 as an accumulation point), one may call the ensuing family of states a *continuous field* whenever the function $\hbar \mapsto \omega_\hbar(\mathcal{Q}_\hbar(f))$ is continuous on $[0, 1]$ for each $f \in C_c^\infty(M)$; this notion is actually intrinsically defined by the continuous field of C^* -algebras, and is therefore independent of the quantization maps \mathcal{Q}_\hbar . In particular, one has

$$(9) \quad \lim_{\hbar \rightarrow 0} \omega_\hbar(\mathcal{Q}_\hbar(f)) = \omega_0(f).$$

Eq. (3) (or (8)) shows that coherent states are indeed examples of continuous fields of states, with the additional property that each ω_\hbar is pure. As an example where all states ω_\hbar are mixed, we mention the convergence of quantum-mechanical

²⁰³Here $\mathcal{S}(\mathbb{R}^{2n})$ is the usual Schwartz space of smooth test functions with rapid decay at infinity.

²⁰⁴For example, coherent states of the type introduced by Perelomov [1986] fit into our setting as follows [Simon, 1980]. Let G be a compact connected Lie group, and \mathcal{O}_λ an integral coadjoint orbit, corresponding to a highest weight λ . (One may think here of $G = SU(2)$ and $\lambda = 0, 1/2, 1, \dots$) Note that $\mathcal{O}_\lambda \cong G/T$, where T is the maximal torus in G with respect to which weights are defined. Let $\mathcal{H}_\lambda^{\text{hw}}$ be the carrier space of the irreducible representation $U_\lambda(G)$ with highest weight λ , containing the highest weight vector Ω_λ . (For $G = SU(2)$ one has $\mathcal{H}_j^{\text{hw}} = \mathbb{C}^{2j+1}$, the well-known Hilbert space of spin j , in which Ω_j is the vector with spin j in the z -direction.) For $\hbar = 1/k$, $k \in \mathbb{N}$, define $\mathcal{H}_\hbar := \mathcal{H}_{\lambda/\hbar}^{\text{hw}}$. Choosing a section $\sigma : \mathcal{O}_\lambda \rightarrow G$ of the projection $G \rightarrow G/T$, one then obtains coherent states $x \mapsto U_{\lambda/\hbar}(\sigma(x))\Omega_{\lambda/\hbar}$ with respect to the Liouville measure on \mathcal{O}_λ and $c_\hbar = \dim(\mathcal{H}_{\lambda/\hbar}^{\text{hw}})$. These states are obviously not defined for all values of \hbar in $(0, 1]$, but only for the discrete set $1/\mathbb{N}$.

²⁰⁵This idea is also confirmed by the fact that at least Schrödinger’s coherent states are states of minimal uncertainty; cf. the references in footnote 129.

²⁰⁶The use of this concept in various mathematical approaches to quantization is basically folklore. For the C^* -algebraic setting see [Emch, 1984; Rieffel, 1989b; Werner, 1995; Blanchard, 1996; Landsman, 1998], and [Nagy, 2000].

partition functions to their classical counterparts in statistical mechanics along these lines; see [Lieb, 1973; Simon, 1980; Duffield, 1990], and [Nourrigat and Royer, 2004]. Finally, one encounters the phenomenon that pure quantum states may coverge to mixed classical ones. The first example of this has just been exhibited in (7); other cases in point are energy eigenstates and WKB states (see Subsections 5.4, 5.5, and 5.6 below).

5.2 Convergence of quantum dynamics to classical motion

Nonrelativistic quantum mechanics is based on the Schrödinger equation (3), which more generally reads

$$(10) \quad H\Psi(t) = i\hbar \frac{\partial\Psi}{\partial t}.$$

The formal solution with initial value $\Psi(0) = \Psi$ is

$$(11) \quad \Psi(t) = e^{-\frac{it}{\hbar}H}\Psi.$$

Here we have assumed that H is a given self-adjoint operator on the Hilbert space \mathcal{H} of the system, so that this solution indeed exists and evolves unitarily by Stone's theorem; cf. [Reed and Simon, 1972] and [Simon, 1976]. Equivalently, one may transfer the time-evolution from states (Schrödinger picture) to operators (Heisenberg picture) by putting

$$(12) \quad A(t) = e^{\frac{it}{\hbar}H}Ae^{-\frac{it}{\hbar}H}.$$

We here restrict ourselves to particle motion in \mathbb{R}^n , so that $\mathcal{H} = L^2(\mathbb{R}^n)$.²⁰⁷ In that case, H is typically given by a formal expression like (3) (on some specific domain).²⁰⁸ Now, the first thing that comes to mind is *Ehrenfest's Theorem* [1927], which states that for any (unit) vector $\Psi \in L^2(\mathbb{R}^n)$ in the domain of $\mathcal{Q}_\hbar(q^j) = x^j$ and $\partial V(x)/\partial x^j$ one has

$$(13) \quad m \frac{d^2}{dt^2} \langle x^j \rangle(t) = - \left\langle \frac{\partial V(x)}{\partial x^j} \right\rangle(t),$$

with the notation

$$(14) \quad \begin{aligned} \langle x^j \rangle(t) &= (\Psi(t), x^j \Psi(t)); \\ \left\langle \frac{\partial V(x)}{\partial x^j} \right\rangle(t) &= \left(\Psi(t), \frac{\partial V(x)}{\partial x^j} \Psi(t) \right). \end{aligned}$$

This looks like Newton's second law for the expectation value of x in the state ψ , with the tiny but crucial difference that Newton would have liked to see $(\partial V/\partial x^j)(\langle x \rangle(t))$ on the right-hand side of (13). Furthermore, even apart from

²⁰⁷See [Hunziker and Sigal, 2000] for a recent survey of N -body Schrödinger operators.

²⁰⁸One then has to prove self-adjointness (or the lack of it) on a larger domain on which the operator is closed; see the literature cited in footnote 42.

this point Ehrenfest's Theorem by no means suffices to have classical behaviour, since it gives no guarantee whatsoever that $\langle x \rangle(t)$ behaves like a point particle. Much of what follows can be seen as an attempt to sharpen Ehrenfest's Theorem to the effect that it *does* indeed yield appropriate classical equations of motion for the expectation values of suitable operators.

We assume that the quantum Hamiltonian has the more general form

$$(15) \quad H = h(\mathcal{Q}_{\hbar}(p_j), \mathcal{Q}_{\hbar}(q^j)),$$

where h is the classical Hamiltonian (i.e. a function defined on classical phase space \mathbb{R}^{2n}) and $\mathcal{Q}_{\hbar}(p_j)$ and $\mathcal{Q}_{\hbar}(q^j)$ are the operators given in (2). Whenever this expression is ambiguous (as in cases like $h(p, q) = pq$), one has to assume a specific quantization prescription such as Weyl quantization \mathcal{Q}_{\hbar}^W (cf. (29)), so that formally one has

$$(16) \quad H = \mathcal{Q}_{\hbar}^W(h).$$

In fact, in the literature to be cited an even larger class of quantum Hamiltonians is treated by the methods explained here. The quantum Hamiltonian H carries an explicit (and rather singular) \hbar -dependence, and for $\hbar \rightarrow 0$ one then expects (11) or (12) to be related in one way or another to the flow of the classical Hamiltonian h . This relationship was already foreseen by Schrödinger (1926a), and was formalized almost immediately after the birth of quantum mechanics by the well-known WKB approximation (cf. [Landau and Lifshitz, 1977] and Subsection 5.5 below). A mathematically rigorous understanding of this and analogous approximation methods only emerged much later, when a technique called *microlocal analysis* was adapted from its original setting of partial differential equations [Hörmander, 1965; Kohn and Nirenberg, 1965; Duistermaat, 1974; 1996; Guillemin and Sternberg, 1977; Howe, 1980; Hörmander, 1979; 1985a; 1985b; Grigis and Sjöstrand, 1994] to the study of the $\hbar \rightarrow 0$ limit of quantum mechanics. This adaptation (often called *semiclassical analysis*) and its results have now been explained in various reviews written by the main players, notably [Robert, 1987; 1998; Helffer, 1988; Paul and Uribe, 1995; Colin de Verdière, 1998; Ivrii, 1998; Dimassi and Sjöstrand, 1999], and [Martinez, 2002] (see also the papers in [Robert, 1992]). More specific references will be given below.²⁰⁹

As mentioned before, the relationship between H and h provided by semiclassical analysis is double-edged. On the one hand, one obtains approximate solutions of (11) or (12), or approximate energy eigenvalues and energy eigenfunctions (sometimes called quasi-modes) for small values of \hbar in terms of classical data. This is how the results are usually presented; one computes specific properties of quantum theory in a certain regime in terms of an underlying classical theory. On the other hand, however, with some effort the very same results can often be reinterpreted as a partial explanation of the emergence of classical dynamics from

²⁰⁹For the heuristic theory of semiclassical asymptotics [Landau and Lifshitz, 1977] is a goldmine.

quantum mechanics. It is the latter aspect of semiclassical analysis, somewhat understated in the literature, that is of interest to us. In this and the next three subsections we restrict ourselves to the simplest type of results, which nonetheless provide a good flavour of what can be achieved and understood by these means. By the same token, we just work with the usual flat phase space $M = \mathbb{R}^{2n}$ as before.

The simplest of all results relating classical and quantum dynamics is this:²¹⁰

If the classical Hamiltonian $h(p, q)$ is at most quadratic in p and q , and the Hamiltonian in (12) is given by (16), then

$$(17) \quad \mathcal{Q}_\hbar^W(f)(t) = \mathcal{Q}_\hbar^W(f_t).$$

Here f_t is the solution of the classical equation of motion $df_t/dt = \{h, f_t\}$; equivalently, one may write

$$(18) \quad f_t(p, q) = f(p(t), q(t)),$$

where $t \mapsto (p(t), q(t))$ is the classical Hamiltonian flow of h with initial condition $(p(0), q(0)) = (p, q)$. This holds for all decent f , e.g., $f \in \mathcal{S}(\mathbb{R}^{2n})$.

This result explains quantum in terms of classical, but the converse may be achieved by combining (17) with (9). This yields

$$(19) \quad \lim_{\hbar \rightarrow 0} \omega_\hbar(\mathcal{Q}_\hbar(f)(t)) = \omega_0(f_t)$$

for any continuous field of states (ω_\hbar) . In particular, for Schrödinger's coherent states (20) one obtains

$$(20) \quad \lim_{\hbar \rightarrow 0} \left(\Psi_{(p,q)}^\hbar, \mathcal{Q}_\hbar(f)(t) \Psi_{(p,q)}^\hbar \right) = f_t(p, q).$$

Now, whereas (17) merely reflects the good symmetry properties of Weyl quantization,²¹¹ (and is false for \mathcal{Q}_\hbar^B), eq. (20) is actually valid for a large class of realistic Hamiltonians and for any deformation quantization map \mathcal{Q}_\hbar that is asymptotically equal to \mathcal{Q}_\hbar^W (cf. (32)). A result of this type was first established by Hepp [1974]; further work in this direction includes [Yajima, 1979; Høegreave *et al.*, 1983; Wang, 1986; Robinson, 1988a; 1988b; Combescure, 1992; Arai, 1995; Combescure and Robert, 1997; Robert, 1998], and [Landsman, 1998].

Impressive results are available also in the Schrödinger picture. The counterpart of (17) is that for any suitably smooth classical Hamiltonian h (even a time-dependent one) that is at most quadratic in the canonical coordinates p and q on

²¹⁰More generally, Egorov's Theorem states that for a large class of Hamiltonians one has $\mathcal{Q}_\hbar^W(f)(t) = \mathcal{Q}_\hbar^W(f_t) + \mathcal{O}(\hbar)$. See, e.g., [Robert, 1987; Dimassi and Sjöstrand, 1999], and [Martinez, 2002].

²¹¹Eq. (17) is equivalent to the covariance of Weyl quantization under the affine symplectic group; cf. footnote 153.

phase space \mathbb{R}^{2n} one may construct generalized coherent states $\Psi_{(p,q,C)}^{\hbar}$, labeled by a set C of classical parameters dictated by the form of h , such that

$$(21) \quad e^{-\frac{it}{\hbar} \mathcal{Q}_h^W(h)} \Psi_{(p,q,C)}^{\hbar} = e^{iS(t)/\hbar} \Psi_{(p(t),q(t),C(t))}^{\hbar}.$$

Here $S(t)$ is the action associated with the classical trajectory $(p(t), q(t))$ determined by h , and $C(t)$ is a solution of a certain system of differential equations that has a classical interpretation as well [Hagedorn, 1998]. Schrödinger's coherent states (20) are a special case for the standard harmonic oscillator Hamiltonian. For more general Hamiltonians one then has an asymptotic result [Hagedorn and Joye, 1999; 2000]²¹²

$$(22) \quad \lim_{\hbar \rightarrow 0} \left\| e^{-\frac{it}{\hbar} \mathcal{Q}_h^W(h)} \Psi_{(p,q,C)}^{\hbar} - e^{iS(t)/\hbar} \Psi_{(p(t),q(t),C(t))}^{\hbar} \right\| = 0.$$

Once again, at first sight such results merely contribute to the understanding of quantum dynamics in terms of classical motion. As mentioned, they may be converted into statements on the emergence of classical motion from quantum mechanics by taking expectation values of suitable \hbar -dependent observables of the type $\mathcal{Q}_h^W(f)$.

For finite \hbar , the second term in (22) is a good approximation to the first - the error even being as small as $\mathcal{O}(\exp(-\gamma/\hbar))$ for some $\gamma > 0$ as $\hbar \rightarrow 0$ - whenever t is smaller than the so-called *Ehrenfest time*

$$(23) \quad T_E = \lambda^{-1} \log(\hbar^{-1}),$$

where λ is a typical inverse time scale of the Hamiltonian (e.g., for chaotic systems it is the largest Lyapunov exponent).²¹³ This is the typical time scale on which semiclassical approximations to wave packet solutions of the time-dependent Schrödinger equation with a general Hamiltonian tend to be valid [Ehrenfest, 1927; Berry *et al.*, 1979; Zaslavsky, 1981; Combescure and Robert, 1997; Bambusi *et al.*, 1999; Hagedorn and Joye, 2000].²¹⁴ For example, Ehrenfest [1927] himself estimated that for a mass of 1 gram a wave packet would double its width only in

²¹²See also [Paul and Uribe, 1995; 1996] as well as the references listed after (20) for analogous statements.

²¹³Recall that throughout this section we assume that \hbar has been made dimensionless through an appropriate rescaling.

²¹⁴One should distinguish here between two distinct approximation methods to the time-dependent Schrödinger equation. Firstly, one has the semiclassical propagation of a quantum-mechanical wave packet, i.e. its propagation as computed from the time-dependence of the parameters on which it depends *according to the underlying classical equations of motion*. It is shown in the references just cited that this approximates the full quantum-mechanical propagation of the wave packet well until $t \sim T_E$. Secondly, one has the time-dependent WKB approximation (for integrable systems) and its generalization to chaotic systems (which typically involve tens of thousands of terms instead of a single one). This second approximation is valid on a much longer time scale, typically $t \sim \hbar^{-1/2}$ [O'Connor, Tomsovic, and Heller, 1992; Heller and Tomsovic, 1993; Tomsovic and Heller, 1993; 2002; Vanicek and Heller, 2003]. Adding to the confusion, Ballentine has claimed over the years that even the semiclassical propagation of a wave packet approximates its quantum-mechanical propagation for times much longer than the Ehrenfest time, typically $t \sim \hbar^{-1/2}$ [Ballentine *et al.*, 1994; Ballentine, 2002; 2003]. This claim

about 10^{13} years under free motion. However, Zurek and Paz [1995] have estimated the Ehrenfest time for Saturn's moon Hyperion to be of the order of 20 years! This obviously poses a serious problem for the program of deriving (the appearance of) classical behaviour from quantum mechanics, which affects all interpretations of this theory.

Finally, we have not discussed the important problem of combining the limit $t \rightarrow \infty$ with the limit $\hbar \rightarrow 0$; this should be done in such a way that T_E is kept fixed. This double limit is of particular importance for quantum chaos; see Robert [1998] and most of the literature cited in Subsection 5.6.

5.3 Wigner functions

The $\hbar \rightarrow 0$ limit of quantum mechanics is often discussed in terms of the so-called *Wigner function*, introduced by Wigner [1932].²¹⁵ Each unit vector (i.e. wave function) $\Psi \in L^2(\mathbb{R}^n)$ defines such a function W_{Ψ}^{\hbar} on classical phase space $M = \mathbb{R}^{2n}$ by demanding that for each $f \in \mathcal{S}(\mathbb{R}^{2n})$ one has

$$(24) \quad (\Psi, \mathcal{Q}_{\hbar}^W(f)\Psi) = \int_{\mathbb{R}^{2n}} \frac{d^n p d^n q}{(2\pi)^n} W_{\Psi}^{\hbar}(p, q) f(p, q).$$

The existence of such a function may be proved by writing it down explicitly as

$$(25) \quad W_{\Psi}^{\hbar}(p, q) = \int_{\mathbb{R}^n} d^n v e^{ipv} \overline{\Psi(q + \frac{1}{2}\hbar v)} \Psi(q - \frac{1}{2}\hbar v).$$

In other words, the quantum-mechanical expectation value of the Weyl quantization of the classical observable f in a quantum state Ψ formally equals the classical expectation value of f with respect to the distribution W_{Ψ} . However, the latter may not be regarded as a probability distribution because it is not necessarily positive definite.²¹⁶ Despite this drawback, the Wigner function possesses some attractive properties. For example, one has

$$(26) \quad \mathcal{Q}_{\hbar}^W(W_{\Psi}^{\hbar}) = \hbar^{-n}[\Psi].$$

This somewhat perverse result means that if the Wigner function defined by Ψ is seen as a classical observable (despite its manifest \hbar -dependence!), then its Weyl

is based on the criterion that the quantum and classical (i.e. Liouville) probabilities are approximately equal on such time scales, but the validity of this criterion hinges on the “statistical” or “ensemble” interpretation of quantum mechanics. According to this interpretation, a pure state provides a description of certain statistical properties of an ensemble of similarly prepared systems, but need not provide a complete description of an individual system. See [Ballentine, 1970; 1986]. Though once defended by von Neumann, Einstein and Popper, this interpretation has now completely fallen out of fashion.

²¹⁵The original context was quantum statistical mechanics; one may write down (24) for mixed states as well. See [Hillery *et al.*, 1984] for a survey.

²¹⁶Indeed, it may not even be in $L^1(\mathbb{R}^{2n})$, so that its total mass is not necessarily defined, let alone equal to 1. Conditions for the positivity of Wigner functions defined by pure states are given by Hudson [1974]; see [Bröcker and Werner, 1995] for the case of mixed states.

quantization is precisely (\hbar^{-n} times) the projection operator onto Ψ .²¹⁷ Furthermore, one may derive the following formula for the transition probability:²¹⁸

$$(27) \quad |(\Phi, \Psi)|^2 = \hbar^n \int_{\mathbb{R}^{2n}} \frac{d^n p d^n q}{(2\pi)^n} W_{\Psi}^{\hbar}(p, q) W_{\Phi}^{\hbar}(p, q).$$

This expression has immediate intuitive appeal, since the integrand on the right-hand side is supported by the area in phase space where the two Wigner functions overlap, which is well in tune with the idea of a transition probability.

The potential lack of positivity of a Wigner function may be remedied by noting that Berezin’s deformation quantization scheme (see (28)) analogously defines functions B_{Ψ}^{\hbar} on phase space by means of

$$(28) \quad (\Psi, \mathcal{Q}_{\hbar}^B(f)\Psi) = \int_{\mathbb{R}^{2n}} \frac{d^n p d^n q}{(2\pi)^n} B_{\Psi}^{\hbar}(p, q) f(p, q).$$

Formally, (27) and (28) immediately yield

$$(29) \quad B_{\Psi}^{\hbar}(p, q) = \hbar^{-n} |(\Psi_{(p,q)}^{\hbar}, \Psi)|^2$$

in terms of Schrödinger’s coherent states (20). This expression is manifestly positive definite. The existence of B_{Ψ}^{\hbar} may be proved rigorously by recalling that the Berezin quantization map $f \mapsto \mathcal{Q}_{\hbar}^B(f)$ is *positive* from $C_0(\mathbb{R}^{2n})$ to $\mathcal{B}(L^2(\mathbb{R}^n))$. This implies that for each (unit) vector $\Psi \in L^2(\mathbb{R}^n)$ the map $f \mapsto (\Psi, \mathcal{Q}_{\hbar}^B(f)\Psi)$ is positive from $C_c(\mathbb{R}^{2n})$ to \mathbb{C} , so that (by the Riesz theorem of measure theory) there must be a measure μ_{Ψ} on \mathbb{R}^{2n} such that $(\Psi, \mathcal{Q}_{\hbar}^B(f)\Psi) = \int d\mu_{\Psi} f$. This measure, then, is precisely given by $d\mu_{\Psi}(p, q) = (2\pi)^{-n} d^n p d^n q B_{\Psi}^{\hbar}(p, q)$. If $(\Psi, \Psi) = 1$, then μ_{Ψ} is a probability measure. Accordingly, despite its \hbar -dependence, B_{Ψ}^{\hbar} defines a bona fide classical probability distribution on phase space, in terms of which one might attempt to visualize quantum mechanics to some extent.

For finite values of \hbar , the Wigner and Berezin distribution functions are different, because the quantization maps \mathcal{Q}_{\hbar}^W and \mathcal{Q}_{\hbar}^B are. The connection between B_{Ψ}^{\hbar} and W_{Ψ}^{\hbar} is easily computed to be

$$(30) \quad B_{\Psi}^{\hbar} = W_{\Psi}^{\hbar} * g^{\hbar},$$

where g^{\hbar} is the Gaussian function

$$(31) \quad g^{\hbar}(p, q) = (2/\hbar)^n \exp(-(p^2 + q^2)/\hbar).$$

This is how physicists look at the Berezin function,²¹⁹ viz. as a Wigner function smeared with a Gaussian so as to become positive. But since g^{\hbar} converges to a

²¹⁷In other words, W_{Ψ} is the Weyl symbol of the projection operator $[\Psi]$.

²¹⁸This formula is well defined since $\Psi \in L^2(\mathbb{R}^n)$ implies $W_{\Psi}^{\hbar} \in L^2(\mathbb{R}^{2n})$.

²¹⁹The ‘Berezin’ functions B_{Ψ}^{\hbar} were introduced by Husimi [1940] from a different point of view, and are therefore actually called *Husimi functions* by physicists.

Dirac delta function as $\hbar \rightarrow 0$ (with respect to the measure $(2\pi)^{-n} d^n p d^n q$ in the sense of distributions), it is clear from (30) that as distributions one has²²⁰

$$(32) \quad \lim_{\hbar \rightarrow 0} (B_{\Psi}^{\hbar} - W_{\Psi}^{\hbar}) = 0.$$

See also (32). Hence in the study of the limit $\hbar \rightarrow 0$ there is little advantage in the use of Wigner functions; quite to the contrary, in limiting procedures their generic lack of positivity makes them more difficult to handle than Berezin functions.²²¹ For example, one would like to write the asymptotic behaviour (8) of coherent states in the form $\lim_{\hbar \rightarrow 0} W_{\Psi_z^{\hbar}}^{\hbar} = \delta_z$. Although this is indeed true in the sense of distributions, the corresponding limit

$$(33) \quad \lim_{\hbar \rightarrow 0} B_{\Psi_z^{\hbar}}^{\hbar} = \delta_z,$$

exists in the sense of (probability) measures, and is therefore defined on a much larger class of test functions.²²² Here and in what follows, we abuse notation: if μ^0 is some probability measure on \mathbb{R}^{2n} and (Ψ^{\hbar}) is a sequence of unit vectors in $L^2(\mathbb{R}^n)$ indexed by \hbar (and perhaps other labels), then $B_{\Psi^{\hbar}}^{\hbar} \rightarrow \mu^0$ for $\hbar \rightarrow 0$ by definition means that for any $f \in C_c^\infty(\mathbb{R}^{2n})$ one has²²³

$$(34) \quad \lim_{\hbar \rightarrow 0} (\Psi^{\hbar}, \mathcal{Q}_{\hbar}^B(f)\Psi^{\hbar}) = \int_{\mathbb{R}^{2n}} d\mu^0 f.$$

5.4 The classical limit of energy eigenstates

Having dealt with coherent states Ψ_z^{\hbar} in (33), in this subsection we discuss the much more difficult problem of computing the limit measure μ^0 for eigenstates Ψ_n^{\hbar} of the quantum Hamiltonian H . Thus we assume that H has eigenvalues E_n^{\hbar} labeled by $n \in \mathbb{N}$ (defined with or without 0 according to convenience), and also depending on \hbar because of the explicit dependence of H on this parameter. The associated eigenstates Ψ_n^{\hbar} then by definition satisfy

$$(35) \quad H\Psi_n^{\hbar} = E_n^{\hbar}\Psi_n^{\hbar}.$$

Here we incorporate the possibility that the eigenvalue E_n^{\hbar} is degenerate, so that the label n extends n . For example, for the one-dimensional harmonic oscillator one has $E_n^{\hbar} = \hbar\omega(n + \frac{1}{2})$ ($n = 0, 1, 2, \dots$) without multiplicity, but for the hydrogen

²²⁰ Eq. (32) should be interpreted as a limit of the distribution on $\mathcal{D}(\mathbb{R}^{2n})$ or $\mathcal{S}(\mathbb{R}^{2n})$ defined by $B_{\Psi}^{\hbar} - W_{\Psi}^{\hbar}$. Both functions are continuous for $\hbar > 0$, but lose this property in the limit $\hbar \rightarrow 0$, generally converging to distributions.

²²¹ See, however, [Robinett, 1993] and [Arai, 1995]. It should be mentioned that (32) expresses the asymptotic equivalence of Wigner and Berezin functions as distributions on \hbar -independent test functions. Even in the limit $\hbar \rightarrow 0$ one is sometimes interested in studying $O(\hbar)$ phenomena, in which case one should make a choice.

²²² Namely those in $C_0(\mathbb{R}^{2n})$ rather than in $\mathcal{D}(\mathbb{R}^{2n})$ or $\mathcal{S}(\mathbb{R}^{2n})$.

²²³ Since \mathcal{Q}_{\hbar}^B may be extended from $C_c^\infty(\mathbb{R}^{2n})$ to $L^\infty(\mathbb{R}^{2n})$, one may omit the stipulation that μ^0 be a probability measure in this definition if one requires convergence for all $f \in L^\infty(\mathbb{R}^{2n})$, or just for all f in the unitization of the C^* -algebra $C_0(\mathbb{R}^{2n})$.

atom the Bohrian eigenvalues $E_n^{\hbar} = -m_e e^4 / 2\hbar^2 n^2$ (where m_e is the mass of the electron and e is its charge) are degenerate, with the well-known eigenfunctions $\Psi_{(n,l,m)}^{\hbar}$ [Landau and Lifshitz, 1977]. Hence in this case one has $\mathfrak{n} = (n, l, m)$ with $n = 1, 2, 3, \dots$, subject to $l = 0, 1, \dots, n - 1$, and $m = -l, \dots, l$.

In any case, it makes sense to let $\mathfrak{n} \rightarrow \infty$; this certainly means $n \rightarrow \infty$, and may in addition involve sending the other labels in \mathfrak{n} to infinity (subject to the appropriate restrictions on $\mathfrak{n} \rightarrow \infty$, as above). One then expects classical behaviour à la Bohr if one simultaneously lets $\hbar \rightarrow 0$ whilst $E_n^{\hbar} \rightarrow E^0$ converges to some ‘classical’ value E^0 . Depending on how one lets the possible other labels behave in this limit, this may also involve similar asymptotic conditions on the eigenvalues of operators commuting with H — see below for details in the integrable case. We denote the collection of such eigenvalues (including E_n^{\hbar}) by $E_{\mathfrak{n}}^{\hbar}$. (Hence in the case where the energy levels E_n^{\hbar} are nondegenerate, the label E is just E .) In general, we denote the collective limit of the eigenvalues $E_{\mathfrak{n}}^{\hbar}$ as $\hbar \rightarrow 0$ and $\mathfrak{n} \rightarrow \infty$ by E^0 .

For example, for the hydrogen atom one has the additional operators J^2 of total angular momentum as well as the operator J_3 of angular momentum in the z -direction. The eigenfunction $\Psi_{(n,l,m)}^{\hbar}$ of H with eigenvalue E_n^{\hbar} is in addition an eigenfunction of J^2 with eigenvalue $j_{\hbar}^2 = \hbar^2 l(l+1)$ and of J_3 with eigenvalue $j_3^{\hbar} = \hbar m$. Along with $n \rightarrow \infty$ and $\hbar \rightarrow 0$, one may then send $l \rightarrow \infty$ and $m \rightarrow \pm\infty$ in such a way that j_{\hbar}^2 and j_3^{\hbar} approach specific constants.

The object of interest, then, is the measure on phase space obtained as the limit of the Berezin functions (29), i.e.

$$(36) \quad \mu_E^0 = \lim_{\hbar \rightarrow 0, \mathfrak{n} \rightarrow \infty} B_{\Psi_{\mathfrak{n}}^{\hbar}}.$$

Although the pioneers of quantum mechanics were undoubtedly interested in quantities like this, it was only in the 1970s that rigorous results were obtained. Two cases are well understood: in this subsection we discuss the *integrable* case, leaving chaotic and more generally *ergodic* motion to Subsection 5.6.

In the physics literature, it was argued that for an integrable system the limiting measure μ_E^0 is concentrated (in the form of a δ -function) on the invariant torus associated to E^0 [Berry, 1977a].²²⁴ Independently, mathematicians began to study a quantity very similar to μ_E^0 , defined by limiting sequences of eigenfunctions of the Laplacian on a Riemannian manifold M . Here the underlying classical flow is Hamiltonian as well, the corresponding trajectories being the geodesics of the given metric (see, for example, [Klingenberg, 1982; Abraham and Marsden, 1985; Katok and Hasselblatt, 1995], or [Landsman, 1998]).²²⁵ The ensuing picture largely confirms the folklore of the physicists:

In the integrable case the limit measure μ_E^0 is concentrated on invariant tori.

²²⁴This conclusion was, in fact, reached from the Wigner function formalism. See [Ozorio de Almeida, 1988] for a review of work of Berry and his collaborators on this subject.

²²⁵The simplest examples of integrable geodesic motion are n -tori, where the geodesics are projections of lines, and the sphere, where the geodesics are great circles [Katok and Hasselblatt, 1995].

See [Charbonnel, 1986; 1988; Zelditch, 1990, 1996a; Toth, 1996; 1999; Nadirashvili *et al.*, 2001], and [Toth and Zelditch, 2002; 2003a; 2003b].²²⁶ Finally, as part of the transformation of microlocal analysis to semiclassical analysis (cf. Subsection 5.2), these results were adapted to quantum mechanics [Paul and Uribe, 1995; 1996].

Let us now give some details for integrable systems (of Liouville type); these include the hydrogen atom as a special case. Integrable systems are defined by the property that on a $2p$ -dimensional phase space M one has p independent²²⁷ classical observables $(f_1 = h, f_2, \dots, f_p)$ whose mutual Poisson brackets all vanish [Arnold, 1989]. One then hopes that an appropriate quantization scheme \mathcal{Q}_\hbar exists under which the corresponding quantum observables $(\mathcal{Q}_\hbar(f_1) = H, \mathcal{Q}_\hbar(f_2), \dots, \mathcal{Q}_\hbar(f_p))$ are all self-adjoint and mutually commute (on a common core).²²⁸ This is indeed the case for the hydrogen atom, where (f_1, f_2, f_3) may be taken to be (h, j^2, j_3) (where j^2 is the total angular momentum and j_3 is its z -component),²²⁹ H is given by (16), $J^2 = \mathcal{Q}_\hbar^W(j^2)$, and $J_3 = \mathcal{Q}_\hbar^W(j_3)$. In general, the energy eigenfunctions Ψ_n^\hbar will be joint eigenfunctions of the operators $(\mathcal{Q}_\hbar(f_1), \dots, \mathcal{Q}_\hbar(f_p))$, so that $E_n^\hbar = (E_{n_1}^\hbar, \dots, E_{n_p}^\hbar)$, with $\mathcal{Q}_\hbar(f_k)\Psi_n^\hbar = E_{n_k}^\hbar \Psi_n^\hbar$. We assume that the submanifolds $\cap_{k=1}^p f_k^{-1}(x_k)$ are compact and connected for each $x \in \mathbb{R}^p$, so that they are tori by the Liouville–Arnold Theorem [Abraham and Marsden, 1985; Arnold, 1989].

Letting $\hbar \rightarrow 0$ and $n \rightarrow \infty$ so that $E_{n_k}^\hbar \rightarrow E_k^0$ for some point $E^0 = (E_1^0, \dots, E_p^0) \in \mathbb{R}^p$, it follows that the limiting measure μ_E^0 as defined in (36) is concentrated on the invariant torus $\cap_{k=1}^p f_k^{-1}(E_k^0)$. This torus is generically p -dimensional, but for singular points E^0 it may be of lower dimension. In particular, in the exceptional circumstance where the invariant torus is one-dimensional, μ_E^0 is concentrated on a classical orbit. Of course, for $p = 1$ (where any Hamiltonian system is integrable) this singular case is generic. Just think of the foliation of \mathbb{R}^2 by the ellipses that form the closed orbits of the harmonic oscillator motion.²³⁰

²²⁶These papers consider the limit $n \rightarrow \infty$ without $\hbar \rightarrow 0$; in fact, a physicist would say that they put $\hbar = 1$. In that case $E_n \rightarrow \infty$; in this procedure the physicists' microscopic $E \sim \mathcal{O}(\hbar)$ and macroscopic $E \sim \mathcal{O}(1)$ regimes correspond to $E \sim \mathcal{O}(1)$ and $E \rightarrow \infty$, respectively.

²²⁷I.e. $df_1 \wedge \dots \wedge df_p \neq 0$ everywhere. At this point we write $2p$ instead of $2n$ for the dimension of phase space in order to avoid notational confusion.

²²⁸There is no general theory of quantum integrable systems. Olshanetsky and Perelomov [1981; 1983] form a good starting point.

²²⁹In fact, if μ is the momentum map for the standard $SO(3)$ -action on \mathbb{R}^3 , then $j^2 = \sum_{k=1}^3 \mu_k^2$ and $j_3 = \mu_3$.

²³⁰It may be enlightening to consider geodesic motion on the sphere; this example may be seen as the hydrogen atom without the radial degree of freedom (so that the degeneracy in question occurs in the hydrogen atom as well). If one sends $l \rightarrow \infty$ and $m \rightarrow \infty$ in the spherical harmonics Y_l^m (which are eigenfunctions of the Laplacian on the sphere) in such a way that $\lim m/l = \cos \varphi$, then the invariant tori are generically two-dimensional, and occur when $\cos \varphi \neq \pm 1$; an invariant torus labeled by such a value of $\varphi \neq 0, \pi$ comprises all great circles (regarded as part of phase space by adding to each point of the geodesic a velocity of unit length and direction tangent to the geodesic) whose angle with the z -axis is φ (more precisely, the angle in question is the one between the normal of the plane through the given great circle and the z -axis). For $\cos \varphi = \pm 1$ (i.e. $m = \pm l$), however, there is only one great circle with $\varphi = 0$ namely the equator (the case

What remains, then, of Bohr's picture of the hydrogen atom in this light?²³¹ Quite a lot, in fact, confirming his remarkable physical intuition. The energy levels Bohr calculated are those given by the Schrödinger equation, and hence remain correct in mature quantum mechanics. His orbits make literal sense only in the "correspondence principle" limit $\hbar \rightarrow 0$, $n \rightarrow \infty$, where, however, the situation is even better than one might expect for integrable systems: because of the high degree of symmetry of the Kepler problem [Guillemin and Sternberg, 1990], one may construct energy eigenfunctions whose limit measure μ^0 concentrates on any desired classical orbit [Nauenberg, 1989].²³² In order to recover a travelling wave packet, one has to form wave packets from a very large number of energy eigenstates with very high quantum numbers, as explained in Subsection 2.4. For finite n and \hbar Bohr's orbits seem to have no meaning, as already recognized by Heisenberg [1969] in his pathfinder days!²³³

5.5 The WKB approximation

One might have expected a section on the $\hbar \rightarrow 0$ limit of quantum mechanics to be centered around the WKB approximation, as practically all textbooks base their discussion of the classical limit on this notion. Although the scope of this method is actually rather limited, it is indeed worth saying a few words about it. For simplicity we restrict ourselves to the time-independent case.²³⁴ In its original formulation, the time-independent WKB method involves an attempt to approximate solutions of the time-independent Schrödinger equation $H\Psi = E\Psi$ by wave functions of the type

$$(37) \quad \Psi(x) = a_{\hbar}(x)e^{\frac{i}{\hbar}S(x)},$$

where a_{\hbar} admits an expansion in \hbar as a power series. Assuming the Hamiltonian H is of the form (15), plugging the Ansatz (37) into the Schrödinger equation, and expanding in \hbar , yields in lowest order the classical (time-independent) Hamilton–Jacobi equation

$$(38) \quad h \left(\frac{\partial S}{\partial x}, x \right) = E,$$

$\varphi = \pi$ corresponds to the same equator traversed in the opposite direction). Hence in this case the invariant torus is one-dimensional. The reader may be surprised that the invariant tori explicitly depend on the choice of variables, but this feature is typical of so-called degenerate systems; see Arnold (1989), §51.

²³¹We ignore coupling to the electromagnetic field here; see footnote 26.

²³²Continuing footnote 230, for a given principal quantum number n one forms the eigenfunction $\Psi_{(n,n-1,n-1)}^{\hbar}$ by multiplying the spherical harmonic Y_{n-1}^{n-1} with the appropriate radial wave function. The limiting measure (36) as $n \rightarrow \infty$ and $\hbar \rightarrow 0$ is then concentrated on an orbit (rather than on an invariant torus). Now, beyond what is possible for general integrable systems, one may use the $SO(4)$ symmetry of the Kepler problem and the construction in footnote 204 for the group-theoretic coherent states of Perelomov [1986] to find the desired eigenfunctions. See also [De Bièvre, 1992] and [De Bièvre *et al.*, 1993].

²³³The later Bohr also conceded this through his idea that causal descriptions are complementary to space-time pictures; see Subsection 3.3.

²³⁴Cf. [Robert, 1998] and references therein for the time-dependent case.

supplemented by the so-called (homogeneous) transport equation²³⁵

$$(39) \quad \left(\frac{1}{2} \Delta S + \sum_k \frac{\partial S}{\partial x^k} \frac{\partial}{\partial x^k} \right) a_0 = 0.$$

In particular, E should be a classically allowed value of the energy. Even when it applies (see below), in most cases of interest the Ansatz (37) is only valid locally (in x), leading to problems with caustics. These problems turn out to be an artefact of the use of the coordinate representation that lies behind the choice of the Hilbert space $\mathcal{H} = L^2(\mathbb{R}^n)$, and can be avoided [Maslov and Fedoriuk, 1981]: the WKB method really comes to its own in a geometric reformulation in terms of symplectic geometry. See [Arnold, 1989; Bates and Weinstein, 1995], and [Dimassi and Sjöstrand, 1999] for (nicely complementary) introductory treatments, and [Guillemin and Sternberg, 1977; Hörmander, 1985a; 1985b], and [Duistermaat, 1974; 1996] for advanced accounts.

The basic observation leading to this reformulation is that in the rare cases that S is defined globally as a smooth function on the configuration space \mathbb{R}^n , it defines a submanifold \mathcal{L} of the phase space $M = \mathbb{R}^{2n}$ by $\mathcal{L} = \{(p = dS(x), q = x), x \in \mathbb{R}^n\}$. This submanifold is *Lagrangian* in having two defining properties: firstly, \mathcal{L} is n -dimensional, and secondly, the restriction of the symplectic form (i.e. $\sum_k dp_k \wedge dq^k$) to \mathcal{L} vanishes. The Hamilton–Jacobi equation (38) then guarantees that the Lagrangian submanifold $\mathcal{L} \subset M$ is contained in the surface $\Sigma_E = h^{-1}(E)$ of constant energy E in M . Consequently, any solution of the Hamiltonian equations of motion that starts in \mathcal{L} remains in \mathcal{L} .

In general, then, the starting point of the WKB approximation is a Lagrangian submanifold $\mathcal{L} \subset \Sigma_E \subset M$, rather than some function S that defines it locally. By a certain adaptation of the geometric quantization procedure, one may, under suitable conditions, associate a unit vector $\Psi_{\mathcal{L}}$ in a suitable Hilbert space to \mathcal{L} , which for small \hbar happens to be a good approximation to an eigenfunction of H at eigenvalue E . This strategy is successful in the integrable case, where the nondegenerate tori (i.e. those of maximal dimension n) provide such Lagrangian submanifolds of M ; the associated unit vector $\Psi_{\mathcal{L}}$ then turns out to be well defined precisely when \mathcal{L} satisfies (generalized) Bohr–Sommerfeld quantization conditions. In fact, this is how the measures μ_E^0 in (36) are generally computed in the integrable case.

If the underlying classical system is not integrable, it may still be close enough to integrability for invariant tori to be defined. Such systems are called quasi-integrable or perturbations of integrable systems, and are described by the Kolmogorov–Arnold–Moser (KAM) theory; see [Gallavotti, 1983; Abraham and Marsden, 1985; Ozorio de Almeida, 1988; Arnold, 1989; Lazutkin, 1993; Gallavotti *et al.*, 2004], and many other books. In such systems the WKB method continues to

²³⁵Only stated here for a classical Hamiltonian $h(p, q) = p^2/2m + V(q)$. Higher-order terms in \hbar yield further, inhomogeneous transport equations for the expansion coefficients $a_j(x)$ in $a_{\hbar} = \sum_j a_j \hbar^j$. These can be solved in a recursive way, starting with (39).

provide approximations to the energy eigenstates relevant to the surviving invariant tori [Colin de Verdière, 1977; Lazutkin, 1993; Popov, 2000], but already loses some of its appeal.

In general systems, notably chaotic ones, the WKB method is almost useless. Indeed, the following theorem of Werner [1995] shows that the measure μ_E^0 defined by a WKB function (37) is concentrated on the Lagrangian submanifold \mathcal{L} defined by S :

Let a_{\hbar} be in $L^2(\mathbb{R}^n)$ for each $\hbar > 0$ with pointwise limit $a_0 = \lim_{\hbar \rightarrow 0} a_{\hbar}$ also in $L^2(\mathbb{R}^n)$,²³⁶ and suppose that S is almost everywhere differentiable. Then for each $f \in C_c^\infty(\mathbb{R}^{2n})$:

$$(40) \quad \lim_{\hbar \rightarrow 0} \left(a_{\hbar} e^{i\frac{S}{\hbar}}, \mathcal{Q}_{\hbar}^B(f) a_{\hbar} e^{i\frac{S}{\hbar}} \right) = \int_{\mathbb{R}^n} d^n x |a_0(x)|^2 f \left(\frac{\partial S}{\partial x}, x \right).$$

As we shall see shortly, this behaviour is impossible for ergodic systems, and this is enough to seal the fate of WKB for chaotic systems in general (except perhaps as a hacker’s tool).

Note, however, that for a given energy level E the discussion so far has been concerned with properties of the classical trajectories on Σ_E (where they are constrained to remain by conservation of energy). Now, it belongs to the essence of quantum mechanics that other parts of phase space than Σ_E might be relevant to the spectral properties of H as well. For example, for a classical Hamiltonian of the simple form $h(p, q) = p^2/2m + V(q)$, this concerns the so-called *classically forbidden area* $\{q \in \mathbb{R}^n \mid V(q) > E\}$ (and any value of p). Here the classical motion can have no properties like integrability or ergodicity, because it does not exist. Nonetheless, and perhaps counterintuitively, it is precisely here that a slight adaptation of the WKB method tends to be most effective. For $q = x$ in the classically forbidden area, the Ansatz (37) should be replaced by

$$(41) \quad \Psi(x) = a_{\hbar}(x) e^{-\frac{S(x)}{\hbar}},$$

where this time S obeys the Hamilton–Jacobi equation ‘for imaginary time’,²³⁷ i.e.

$$(42) \quad h \left(i \frac{\partial S}{\partial x}, x \right) = E,$$

and the transport equation (39) is unchanged. For example, it follows that in one dimension (with a Hamiltonian of the type (3)) the WKB function (41) assumes the form

$$(43) \quad \Psi(x) \sim e^{-\frac{\sqrt{2m}}{\hbar} \int^{|x|} dy \sqrt{V(y) - E}}$$

²³⁶This assumption is not made in Werner (1995), who directly assumes that $\Psi = a_0 \exp(iS/\hbar)$ in (37).

²³⁷This terminology comes from the Lagrangian formalism, where the classical action $S = \int dt L(t)$ is replaced by iS through the substitution $t = -i\tau$ with $\tau \in \mathbb{R}$.

in the forbidden region, which explains both the tunnel effect in quantum mechanics (i.e. the propagation of the wave function into the forbidden region) *and* the fact that this effect disappears in the limit $\hbar \rightarrow 0$. However, even here the use of WKB methods has now largely been superseded by techniques developed by Agmon [1982]; see, for example, [Hislop and Sigal, 1996] and [Dimassi and Sjöstrand, 1999] for reviews.

5.6 *Epilogue: quantum chaos*

Chaos in classical mechanics was probably known to Newton and was famously highlighted by Poincaré (1892–1899),²³⁸ but its relevance for (and potential threat to) quantum theory was apparently first recognized by Einstein [1917] in a paper that was ‘completely ignored for 40 years’ [Gutzwiller, 1992].²³⁹ Currently, the study of quantum chaos is one of the most thriving businesses in all of physics, as exemplified by innumerable conference proceedings and monographs on the subject, ranging from the classic by Gutzwiller [1990] to the online *opus magnum* by Cvitanovic *et al.* [2005].²⁴⁰ Nonetheless, the subject is still not completely understood, and provides a fascinating testing ground for the interplay between classical and quantum mechanics.

One should distinguish between various different goals in the field of quantum chaos. The majority of papers and books on quantum chaos is concerned with the semiclassical analysis of some concretely given quantum system having a chaotic system as its classical limit. This means that one tries to approximate (for small \hbar) a suitable quantum-mechanical expression in terms of data associated with the underlying classical motion. Michael Berry even described this goal as the “Holy Grail” of quantum chaos. The methods described in Subsection 5.2 contribute to this goal, but are largely independent of the nature of the dynamics. In this subsection we therefore concentrate on techniques and results specific to chaotic motion.

Historically, the first new tool in semiclassical approximation theory that specifically applied to chaotic systems was the so-called *Gutzwiller trace formula*.²⁴¹ Roughly speaking, this formula approximates the eigenvalues of the quantum Hamiltonian in terms of the periodic (i.e. closed) orbits of the underlying clas-

²³⁸See also [Diacu and Holmes, 1996] and [Barrow-Green, 1997] for historical background.

²³⁹It was the study of the very same Helium atom that led Heisenberg to believe that a fundamentally new ‘quantum’ mechanics was needed to replace the inadequate old quantum theory of Bohr and Sommerfeld. See [Mehra and Rechenberg, 1982b] and [Cassidy, 1992]. Another microscopic example of a chaotic system is the hydrogen atom in an external magnetic field.

²⁴⁰Other respectable books include, for example, [Guhr *et al.*, 1998, [Haake, 2001] and [Reichl, 2004].

²⁴¹This attribution is based on Gutzwiller [1971]. A similar result was independently derived by Balian and Bloch [1972; 1974]. See also [Gutzwiller, 1990] and [Brack and Bhaduri, 2003] for mathematically heuristic but otherwise excellent accounts of semiclassical physics based on the trace formula. Mathematically rigorous discussions and proofs may be found in [Colin de Verdière, 1973; Duistermaat and Guillemin, 1975; Guillemin and Uribe, 1989; Paul and Uribe, 1995], and [Combes *et al.*, 1999].

sical Hamiltonian.²⁴² The Gutzwiller trace formula does not start from the wave function (as the WKB approximation does), but from the *propagator* $K(x, y, t)$. Physicists write this as $K(x, y, t) = \langle x | \exp(-itH/\hbar) | y \rangle$, whereas mathematicians see it as the Green's function in the formula

$$(44) \quad e^{-\frac{it}{\hbar}H}\Psi(x) = \int d^n y K(x, y, t)\Psi(y),$$

where $\Psi \in L^2(\mathbb{R}^n)$. Its (distributional) Laplace transform

$$(45) \quad G(x, y, E) = \frac{1}{i\hbar} \int_0^\infty dt K(x, y, t) e^{\frac{itE}{\hbar}}$$

contains information about both the spectrum and the eigenfunctions; for if the former is discrete, one has

$$(46) \quad G(x, y, E) = \sum_j \frac{\Psi_j(x)\overline{\Psi_j(y)}}{E - E_j}.$$

It is possible to approximate K or G itself by an expression of the type

$$(47) \quad K(x, y, t) \sim (2\pi i\hbar)^{-n/2} \sum_P \sqrt{|\det V_P|} e^{\frac{i}{\hbar}S_P(x,y,t) - \frac{1}{2}i\pi\mu_P},$$

where the sum is over *all* classical paths P from y to x in time t (i.e. paths that solve the classical equations of motion). Such a path has an associated action S_P , Maslov index μ_P , and Van Vleck [1928] determinant $\det V_P$ [Arnold, 1989]. For chaotic systems one typically has to include tens of thousands of paths in the sum, but if one does so the ensuing approximation turns out to be remarkably successful [Heller and Tomsovic, 1993; Tomsovic and Heller, 1993]. The Gutzwiller trace formula is a semiclassical approximation to

$$(48) \quad g(E) = \int d^n x G(x, x, E) = \sum_j \frac{1}{E - E_j},$$

for a quantum Hamiltonian with discrete spectrum and underlying classical Hamiltonian having chaotic motion. It has the form

$$(49) \quad g(E) \sim g_0(E) + \frac{1}{i\hbar} \sum_P \sum_{k=1}^\infty \frac{T_P}{2 \sinh(k\chi_P/2)} e^{\frac{ik}{\hbar}S_P(E) - \frac{1}{2}i\pi\mu_P},$$

where g_0 is a smooth function giving the mean density of states. This time, the sum is over all (prime) *periodic* paths P of the classical Hamiltonian at energy E with associated action $S_P(E) = \oint pdq$ (where the momentum p is determined by P , given E), period T_P , and stability exponent χ_P (this is a measure of how rapidly neighbouring trajectories drift away from P). Since the frustration expressed by

²⁴²Such orbits are dense but of Liouville measure zero in chaotic classical systems. Their crucial role was first recognized by Poincaré (1892–1899).

Einstein [1917], this was the first indication that semiclassical approximations had some bearing on chaotic systems.

Another important development concerning energy levels was the formulation of two key conjectures:²⁴³

- If the classical dynamics defined by the classical Hamiltonian h is integrable, then the spectrum of H is “uncorrelated” or “random” [Berry and Tabor, 1977].
- If the classical dynamics defined by h is chaotic, then the spectrum of H is “correlated” or “regular” [Bohigas, Giannoni, and Schmit, 1984].

The notions of correlation and randomness used here can be made precise using notions like the distribution of level spacings and the pair correlation function of eigenvalues; see [Zelditch, 1996a] and [De Bièvre, 2001] for introductory treatments, and most of the literature cited in this subsection for further details.²⁴⁴

We now consider energy eigenfunctions instead of eigenvalues, and return to the limit measure (36). In the non (quasi-) integrable case, the key result is that

*for ergodic classical motion,²⁴⁵ the limit measure μ_E^0 coincides with the (normalized) Liouville measure induced on the constant energy surface $\Sigma_E \equiv h^{-1}(E)$.*²⁴⁶

This result was first suggested in the mathematical literature for ergodic geodesic motion on compact hyperbolic Riemannian manifolds [Snirelman, 1974], where it was subsequently proved with increasing generality [Colin de Verdière, 1985; Zelditch, 1987].²⁴⁷ For certain other ergodic systems this property was proved by Zelditch [1991], Gérard and Leichtnam [1993], Zelditch and Zworski [1996], and

²⁴³Strictly speaking, both conjectures are wrong; for example, the harmonic oscillator yields a counterexamples to the first one. See [Zelditch, 1996a] for further information. Nonetheless, the conjectures are believed to be true in a deeper sense.

²⁴⁴This aspect of quantum chaos has applications to number theory and might even lead to a proof of the Riemann hypothesis; see, for example, [Sarnak, 1999; Berry and Keating, 1999], and many other recent papers. Another relevant connection, related to the one just mentioned, is between energy levels and random matrices; see especially [Guhr *et al.*, 1998]. For the plain relevance of all this to practical physics see [Mirlin, 2000].

²⁴⁵Ergodicity is the weakest property that any chaotic dynamical system possesses. See [Katok and Hasselblatt, 1995; Emch and Liu, 2002; Gallavotti *et al.*, 2004], and countless other books.

²⁴⁶The unnormalized Liouville measure μ_E^u on Σ_E is defined by $\mu_E^u(B) = \int_B dS_E(x) (\|dh(x)\|)^{-1}$, where dS_E is the surface element on Σ_E and $B \subset \Sigma_E$ is a Borel set. If Σ_E is compact, the normalized Liouville measure μ_E on Σ_E is given by $\mu_E(B) = \mu_E^u(B)/\mu_E^u(\Sigma_E)$. It is a probability measure on Σ_E , reflecting the fact that the eigenvectors Ψ_n^h are normalized to unit length so as to define quantum-mechanical states.

²⁴⁷In the Riemannian case with $\hbar = 1$ the cosphere bundle S^*Q (i.e. the subbundle of the cotangent bundle T^*Q consisting of one-forms of unit length) plays the role of Σ_E . Low-dimensional examples of ergodic geodesic motion are provided by compact hyperbolic spaces. Also cf. [Zelditch, 1992a] for the physically important case of a particle moving in an external gauge field. See also the appendix to Lazutkin [1993] by A.I. Shnirelman, and [Nadirashvili *et al.*, 2001] for reviews.

others; to the best of our knowledge a completely general proof remains to be given.

An analogous version for Schrödinger operators on \mathbb{R}^n was independently stated in the physics literature [Berry, 1977b; Voros, 1979], and was eventually proved under certain assumptions on the potential by Helffer, Martinez and Robert [1987], Charbonnel [1992], and Paul and Uribe [1995]. Under suitable assumptions one therefore has

$$(50) \quad \lim_{\hbar \rightarrow 0, n \rightarrow \infty} (\Psi_n^\hbar, \mathcal{Q}_\hbar^B(f)\Psi_n^\hbar) = \int_{\Sigma_E} d\mu_E f$$

for any $f \in C_c^\infty(\mathbb{R}^{2n})$, where again μ_E is the (normalized) Liouville measure on $\Sigma_E \subset \mathbb{R}^{2n}$ (assuming this space to be compact). In particular, in the ergodic case μ_E^0 only depends on E^0 and is the same for (almost) every sequence of energy eigenfunctions (Ψ_n^\hbar) as long as $E_n^\hbar \rightarrow E^0$.²⁴⁸ Thus the support of the limiting measure is uniformly spread out over the largest part of phase space that is dynamically possible.

The result that for ergodic classical motion μ_E^0 is the Liouville measure on Σ_E under the stated condition leaves room for the phenomenon of ‘scars’, according to which in chaotic systems the limiting measure is sometimes concentrated on periodic classical orbits. This terminology is used in two somewhat different ways in the literature. ‘Strong’ scars survive in the limit $\hbar \rightarrow 0$ and concentrate on stable closed orbits;²⁴⁹ they may come from ‘exceptional’ sequences of eigenfunctions.²⁵⁰ These are mainly considered in the mathematical literature; cf. [Nadirashvili *et al.*, 2001] and references therein.

In the physics literature, on the other hand, the notion of a scar usually refers to an anomalous concentration of the functions $B_{\Psi_n^\hbar}^\hbar$ (cf. (29)) near *unstable* closed orbits for *finite* values of \hbar ; see [Heller and Tomsovic, 1993; Tomsovic and Heller, 1993; Kaplan and Heller, 1998a; 1998b], and [Kaplan, 1999] for surveys. Such scars turn out to be crucial in attempts to explain the energy spectrum of the associated quantum system. The reason why such scars do not survive the (double) limit in (36) is that this limit is defined with respect to \hbar -independent smooth test functions. Physically, this means that one averages over more and more De Broglie wavelengths as $\hbar \rightarrow 0$, eventually losing information about the single wavelength scale [Kaplan, 1999]. Hence to pick them up in a mathematically sound way, one should redefine (36) as a pointwise limit [Duclos and Høegre, 1993; Paul and Uribe, 1996; 1998]. In any case, there is no contradiction between the mathematical results cited and what physicists have found.

Another goal of quantum chaos is the identification of chaotic phenomena within a given quantum-mechanical model. Here the slight complication arises that one

²⁴⁸ The result is not necessarily valid for all sequences (Ψ_n^\hbar) with the given limiting behaviour, but only for ‘almost all’ such sequences (technically, for a class of sequences of density 1). See, for example, [De Bièvre, 2001] for a simple explanation of this.

²⁴⁹ An orbit $\gamma \subset M$ is called *stable* when for each neighbourhood U of γ there is neighbourhood $V \subset U$ of γ such that $z(t) \in U$ for all $z \in V$ and all t .

²⁵⁰ Cf. footnote 248.

cannot simply copy the classical definition of chaos in terms of diverging trajectories in phase space, since (by unitarity of time-evolution) in quantum mechanics $\|\Psi(t) - \Phi(t)\|$ is constant in time t for solutions of the Schrödinger equation. However, this just indicates that should intrinsic quantum chaos exist, it has to be defined differently from classical chaos.²⁵¹ This has now been largely accomplished in the algebraic formulation of quantum theory [Benatti, 1993; Emch *et al.*, 1994; Zelditch, 1996b; 1996c; Belot and Earman, 1997; Alicki and Fannes, 2001; Narnhofer, 2001]. The most significant recent development in this direction in the “heuristic” literature has been the study of the quantity

$$(51) \quad M(t) = |(e^{-\frac{it}{\hbar}(H+\Sigma)}\Psi, e^{-\frac{it}{\hbar}H}\Psi)|^2,$$

where Ψ is a coherent state (or Gaussian wave packet), and Σ is some perturbation of the Hamiltonian H [Peres, 1984]. In what is generally regarded as a breakthrough in the field, Jalabert and Pastawski [2001] discovered that in a certain regime $M(t)$ is independent of the detailed form of Σ and decays as $\sim \exp(-\lambda t)$, where λ is the (largest) Lyapunov exponent of the underlying classical system. See [Cucchietti, 2004] for a detailed account and further development.

In any case, the possibility that classical chaos appears in the $\hbar \rightarrow 0$ limit of quantum mechanics is by no means predicated on the existence of intrinsic quantum chaos in the above sense.²⁵² For even in the unlikely case that quantum dynamics would turn out to be intrinsically non-chaotic, its classical limit is sufficiently singular to admit kinds of classical motion without a qualitative counterpart in quantum theory. This possibility is not only confirmed by most of the literature on quantum chaos (little of which makes any use of notions of intrinsic quantum chaotic motion), but even more so by the possibility of *incomplete* motion. This is a type of dynamics in which the flow of the Hamiltonian vector field is only defined until a certain time $t_f < \infty$ (or from an initial time $t_i > -\infty$), which means that the equations of motion have no solution for $t > t_f$ (or $t < t_i$).²⁵³ The

²⁵¹As pointed out by Belot and Earman [1997], the Koopman formulation of classical mechanics (cf. footnote 172) excludes classical chaos if this is formulated in terms of trajectories in Hilbert space. The transition from classical to quantum notions of chaos can be smoothened by first reformulating the classical definition of chaos (normally put in terms of properties of trajectories in phase space).

²⁵²Arguments by [Ford, 1988] and others to the effect that quantum mechanics is wrong because it cannot give rise to chaos in its classical limit have to be discarded for the reasons given here. See also [Belot and Earman, 1997]. In fact, using the same argument, such authors could simultaneously have ‘proved’ the *opposite* statement that any classical dynamics that arises as the classical limit of a quantum theory with non-degenerate spectrum must be ergodic. For the naive definition of quantum ergodic flow clearly is that quantum time-evolution sweeps out all states at some energy E ; but for non-degenerate spectra this is a tautology by definition of an eigenfunction!

²⁵³The simplest examples are incomplete Riemannian manifolds Q with geodesic flow; within this class, the case $Q = (0, 1)$ with flat metric is hard to match in simplicity. Clearly, the particle reaches one of the two boundary points in finite time, and does not know what to do (or even whether it exists) afterwards. Other examples come from potentials V on $Q = \mathbb{R}^n$ with the property that the classical dynamics is incomplete; see [Reed and Simon, 1975] and [Gallavotti, 1983]. On a somewhat different note, the Universe itself has incomplete dynamics because of the Big Bang and possible Big Crunch.

point, then, is that unitary quantum dynamics, though intrinsically complete, may very well have incomplete motion as its classical limit.²⁵⁴

6 THE LIMIT $N \rightarrow \infty$

In this section we show to what extent classical physics may approximately emerge from quantum theory when the size of a system becomes large. Strictly classical behaviour would be an idealization reserved for the limit where this size is infinite, which we symbolically denote by “ $\lim N \rightarrow \infty$ ”. As we shall see, mathematically speaking this limit is a special case of the limit $\hbar \rightarrow 0$ discussed in the previous chapter. What is more, we shall show that formally the limit $N \rightarrow \infty$ even falls under the heading of continuous fields of C^* -algebras and deformation quantization (see Subsection 4.3.) Thus the ‘philosophical’ nature of the idealization involved in assuming that a system is infinite is much the same as that of assuming $\hbar \rightarrow 0$ in a quantum system of given (finite) size; in particular, the introductory comments in Section 1 apply here as well.

An analogous discussion pertains to the derivation of thermodynamics from statistical mechanics [Emch and Liu, 2002; Batterman, 2005]. For example, *in theory* phase transitions only occur in infinite systems, but *in practice* one sees them every day. Thus it appears to be valid to approximate a pot of 10^{23} boiling water molecules by an infinite number of such molecules. The basic point is that the distinction between microscopic and macroscopic regimes is unsharp unless one admits infinite systems as an idealization, so that one can simply say that

²⁵⁴ The quantization of the Universe is unknown at present, but geodesic motion on Riemannian manifolds, complete or not, is quantized by $H = -\frac{\hbar^2}{2m}\Delta$ (perhaps with an additional term proportional to the Ricci scalar R , see [Landsman, 1998]), where Δ is the Laplacian, and quantization on $Q = \mathbb{R}^n$ is given by the Schrödinger equation (3), whether or not the classical dynamics is complete. In these two cases, and probably more generally, the incompleteness of the classical motion is often (but not always) reflected by the lack of essential self-adjointness of the quantum Hamiltonian on its natural initial domain $C_c^\infty(Q)$. For example, if Q is complete as a Riemannian manifold, then Δ is essentially self-adjoint on $C_c^\infty(Q)$ [Chernoff, 1973; Strichartz, 1983], and if Q is incomplete then the Laplacian usually fails to be essentially self-adjoint on this domain (but see [Horowitz and Marolf, 1995] for counterexamples). One may refer to the latter property as quantum-mechanical incompleteness [Reed and Simon, 1975], although a Hamiltonian that fails to be essentially self-adjoint on $C_c^\infty(Q)$ can often be extended (necessarily in a non-unique way) to a self-adjoint operator by a choice of boundary conditions (possibly at infinity). By Stone’s theorem, the quantum dynamics defined by each self-adjoint extension is unitary (and therefore defined for all times). Similarly, although no general statement can be made relating (in)complete classical motion in a potential to (lack of) essential selfadjointness of the corresponding Schrödinger operator, it is usually the case that completeness implies essential selfadjointness, and vice versa. See [Reed and Simon, 1975], Appendix to §X.1, where the reader may also find examples of classically incomplete but quantum-mechanically complete motion, and vice versa. Now, here is the central point for the present discussion: as probably first noted by Hepp [1974], *different self-adjoint extensions have the same classical limit* (in the sense of (20) or similar criteria), namely the given *incomplete* classical dynamics. This proves that complete quantum dynamics can have incomplete motion as its classical limit. However, much remains to be understood in this area. See also [Earman, 2005; 2006].

microscopic systems are finite, whereas macroscopic systems are infinite. This procedure is eventually justified by the results it produces.

Similarly, in the context of quantum theory classical behaviour is simply not found in finite systems (when $\hbar > 0$ is fixed), whereas, as we shall see, it *is* found in infinite ones. Given the observed classical nature of the macroscopic world,²⁵⁵ at the end of the day one concludes that the idealization in question is apparently a valid one. One should not be confused by the fact that the error in the number of particles this approximation involves (viz. $\infty - 10^{23} = \infty$) is considerably larger than the number of particles in the actual system. If all of the 10^{23} particles in question were *individually* tracked down, the approximation is indeed a worthless one, but the point is rather that the limit $N \rightarrow \infty$ is valid whenever *averaging* over $N = 10^{23}$ particles is well approximated by averaging over an arbitrarily larger number N (which, then, one might as well let go to infinity). Below we shall give a precise version of this argument.

Despite our opening comments above, the quantum theory of infinite systems has features of its own that deserve a separate section. Our treatment is complementary to texts such as Thirring [1983], Strocchi [1985], Bratteli and Robinson [1987], Haag [1992], Araki [1999], and Sewell [1986; 2002], which should be consulted for further information on infinite quantum systems. The theory in Subsections 6.1 and 6.5 is a reformulation in terms of continuous field of C^* -algebras and deformation quantization of the more elementary parts of a remarkable series of papers on so-called quantum mean-field systems by Raggio and Werner [1989; 1991], Duffield and Werner [1992a,b,c], and Duffield, Roos, and Werner [1992]. These models have their origin in the treatment of the BCS theory of superconductivity due to Bogoliubov [1958] and Haag [1962], with important further contributions by Thirring and Wehrl [1967], Thirring [1968], Hepp [1972], Hepp and Lieb [1973], Rieckers [1984], Morchio and Strocchi [1987], Duffner and Rieckers [1988], Bona [1988; 1989; 2000], Unnerstall [1990a; 1990b], Bagarello and Morchio [1992], Sewell [2002], and others.

6.1 Macroscopic observables

The large quantum systems we are going to study consist of N copies of a single quantum system with unital algebra of observables \mathcal{A}_1 . Almost all features already emerge in the simplest example $\mathcal{A}_1 = M_2(\mathbb{C})$ (i.e. the complex 2×2 matrices), so there is nothing wrong with having this case in mind as abstraction increases.²⁵⁶ The aim of what follows is to describe in what precise sense macroscopic observables (i.e. those obtained by averaging over an infinite number of sites) are “classical”.

²⁵⁵With the well-known mesoscopic exceptions [Leggett, 2002; Brezger *et al.*, 2002; Chiorescu *et al.*, 2003; Marshall *et al.*, 2003; Devoret *et al.*, 2004].

²⁵⁶In the opposite direction of greater generality, it is worth noting that the setting below actually incorporates quantum systems defined on general lattices in \mathbb{R}^n (such as \mathbb{Z}^n). For one could relabel things so as to make $\mathcal{A}_{1/N}$ below the algebra of observables of all lattice points Λ contained in, say, a sphere of radius N . The limit $N \rightarrow \infty$ then corresponds to the limit $\Lambda \rightarrow \mathbb{Z}^n$.

From the single C^* -algebra \mathcal{A}_1 , we construct a continuous field of C^* -algebras $\mathcal{A}^{(c)}$ over

$$(1) \quad I = 0 \cup 1/\mathbb{N} = \{0, \dots, 1/N, \dots, \frac{1}{3}, \frac{1}{2}, 1\} \subset [0, 1],$$

as follows. We put

$$\begin{aligned} \mathcal{A}_0^{(c)} &= C(\mathcal{S}(\mathcal{A}_1)); \\ (2) \quad \mathcal{A}_{1/N}^{(c)} &= \mathcal{A}_1^N, \end{aligned}$$

where $\mathcal{S}(\mathcal{A}_1)$ is the state space of \mathcal{A}_1 (equipped with the weak*-topology)²⁵⁷ and $\mathcal{A}_1^N = \hat{\otimes}^N \mathcal{A}_1$ is the (spatial) tensor product of N copies of \mathcal{A}_1 .²⁵⁸ This explains the suffix c in $\mathcal{A}^{(c)}$: it refers to the fact that the limit algebra $\mathcal{A}_0^{(c)}$ is classical or commutative.

For example, take $\mathcal{A}_1 = M_2(\mathbb{C})$. Each state is given by a density matrix, which is of the form

$$(3) \quad \rho(x, y, z) = \frac{1}{2} \begin{pmatrix} 1+z & x-iy \\ x+iy & 1-z \end{pmatrix},$$

for some $(x, y, z) \in \mathbb{R}^3$ satisfying $x^2 + y^2 + z^2 \leq 1$. Hence $\mathcal{S}(M_2(\mathbb{C}))$ is isomorphic (as a compact convex set) to the three-ball B^3 in \mathbb{R}^3 . The pure states are precisely the points on the boundary,²⁵⁹ i.e. the density matrices for which $x^2 + y^2 + z^2 = 1$ (for these and these alone define one-dimensional projections).²⁶⁰

In order to define the continuous sections of the field, we introduce the *symmetrization maps* $j_{NM} : \mathcal{A}_1^M \rightarrow \mathcal{A}_1^N$, defined by

$$(4) \quad j_{NM}(A_M) = S_N(A_M \otimes 1 \otimes \dots \otimes 1),$$

where one has $N - M$ copies of the unit $1 \in \mathcal{A}_1$ so as to obtain an element of \mathcal{A}_1^N . The symmetrization operator $S_N : \mathcal{A}_1^N \rightarrow \mathcal{A}_1^N$ is given by (linear and continuous)

²⁵⁷In this topology one has $\omega_\lambda \rightarrow \omega$ when $\omega_\lambda(A) \rightarrow \omega(A)$ for each $A \in \mathcal{A}_1$.

²⁵⁸When \mathcal{A}_1 is finite-dimensional the tensor product is unique. In general, one needs the *projective* tensor product at this point. See footnote 90. The point is the same here: any tensor product state $\omega_1 \otimes \dots \otimes \omega_N$ on $\hat{\otimes}^N \mathcal{A}_1$ — defined on elementary tensors by $\omega_1 \otimes \dots \otimes \omega_N(A_1 \otimes \dots \otimes A_N) = \omega_1(A_1) \dots \omega_N(A_N)$ — extends to a state on $\hat{\otimes}^N \mathcal{A}_1$ by continuity.

²⁵⁹The *extreme boundary* $\partial_e K$ of a convex set K consists of all $\omega \in K$ for which $\omega = p\rho + (1-p)\sigma$ for some $p \in (0, 1)$ and $\rho, \sigma \in K$ implies $\rho = \sigma = \omega$. If $K = \mathcal{S}(\mathcal{A})$ is the state space of a C^* -algebra \mathcal{A} , the extreme boundary consists of the pure states on \mathcal{A} (the remainder of $\mathcal{S}(\mathcal{A})$ consisting of mixed states). If K is embedded in a vector space, the extreme boundary $\partial_e K$ may or may not coincide with the geometric boundary ∂K of K . In the case $K = B^3 \subset \mathbb{R}^3$ it does, but for an equilateral triangle in \mathbb{R}^2 it does not, since $\partial_e K$ merely consists of the corners of the triangle whereas the geometric boundary includes the sides as well.

²⁶⁰Eq. (3) has the form $\rho(x, y, z) = \frac{1}{2}(x\sigma_x + y\sigma_y + z\sigma_z)$, where the σ_i are the Pauli matrices. This yields an isomorphism between \mathbb{R}^3 and the Lie algebra of $SO(3)$ in its spin- $\frac{1}{2}$ representation $\mathcal{D}_{1/2}$ on \mathbb{C}^2 . This isomorphism intertwines the defining action of $SO(3)$ on \mathbb{R}^3 with its adjoint action on $M_2(\mathbb{C})$. I.e., for any rotation R one has $\rho(R\mathbf{x}) = \mathcal{D}_{1/2}(R)\rho(\mathbf{x})\mathcal{D}_{1/2}(R)^{-1}$. This will be used later on (see Subsection 6.5).

extension of

$$(5) \quad S_N(B_1 \otimes \cdots \otimes B_N) = \frac{1}{N!} \sum_{\sigma \in \mathfrak{S}_N} B_{\sigma(1)} \otimes \cdots \otimes B_{\sigma(N)},$$

where \mathfrak{S}_N is the permutation group (i.e. symmetric group) on N elements and $B_i \in \mathcal{A}_1$ for all $i = 1, \dots, N$. For example, $j_{N1} : \mathcal{A}_1 \rightarrow \mathcal{A}_1^N$ is given by

$$(6) \quad j_{N1}(B) = \overline{B}^{(N)} = \frac{1}{N} \sum_{k=1}^N 1 \otimes \cdots \otimes B_{(k)} \otimes 1 \cdots \otimes 1,$$

where $B_{(k)}$ is B seen as an element of the k 'th copy of \mathcal{A}_1 in \mathcal{A}_1^N . As our notation $\overline{B}^{(N)}$ indicates, this is just the ‘average’ of B over all copies of \mathcal{A}_1 . More generally, in forming $j_{NM}(A_M)$ an operator $A_M \in \mathcal{A}_1^M$ that involves M sites is averaged over $N \geq M$ sites. When $N \rightarrow \infty$ this means that one forms a *macroscopic* average of an M -particle operator.

We say that a sequence $A = (A_1, A_2, \dots)$ with $A_N \in \mathcal{A}_1^N$ is *symmetric* when

$$(7) \quad A_N = j_{NM}(A_M)$$

for some fixed M and all $N \geq M$. In other words, the tail of a symmetric sequence entirely consists of ‘averaged’ or ‘intensive’ observables, which become macroscopic in the limit $N \rightarrow \infty$. Such sequences have the important property that they commute in this limit; more precisely, if A and A' are symmetric sequences, then

$$(8) \quad \lim_{N \rightarrow \infty} \|A_N A'_N - A'_N A_N\| = 0.$$

As an enlightening special case we take $A_N = j_{N1}(B)$ and $A'_N = j_{N1}(C)$ with $B, C \in \mathcal{A}_1$. One immediately obtains from the relation $[B_{(k)}, C_{(l)}] = 0$ for $k \neq l$ that

$$(9) \quad [\overline{B}^{(N)}, \overline{C}^{(N)}] = \frac{1}{N} \overline{[B, C]}^{(N)}.$$

For example, if $\mathcal{A}_1 = M_2(\mathbb{C})$ and if for B and C one takes the spin- $\frac{1}{2}$ operators $S_j = \frac{\hbar}{2} \sigma_j$ for $j = 1, 2, 3$ (where σ_j are the Pauli matrices), then

$$(10) \quad [\overline{S}_j^{(N)}, \overline{S}_k^{(N)}] = i \frac{\hbar}{N} \epsilon_{jkl} \overline{S}_l^{(N)}.$$

This shows that averaging one-particle operators leads to commutation relations formally like those of the one-particle operators in question, but with Planck’s constant \hbar replaced by a variable \hbar/N . For constant $\hbar = 1$ this leads to the interval (1) over which our continuous field of C^* -algebras is defined; for any other constant value of \hbar the field would be defined over $I = 0 \cup \hbar/\mathbb{N}$, which of course merely changes the labeling of the C^* -algebras in question.

We return to the general case, and denote a section of the field with fibers (2) by a sequence $A = (A_0, A_1, A_2, \dots)$, with $A_0 \in \mathcal{A}_0^{(c)}$ and $A_N \in \mathcal{A}_1^N$ as before

(i.e. the corresponding section is $0 \mapsto A_0$ and $1/N \mapsto A_N$). We then complete the definition of our continuous field by declaring that a sequence A defines a *continuous* section iff:

- (A_1, A_2, \dots) is *approximately symmetric*, in the sense that for any $\varepsilon > 0$ there is an N_ε and a symmetric sequence A' such that $\|A_N - A'_N\| < \varepsilon$ for all $N \geq N_\varepsilon$,²⁶¹
- $A_0(\omega) = \lim_{N \rightarrow \infty} \omega^N(A_N)$, where $\omega \in \mathcal{S}(\mathcal{A}_1)$ and $\omega^N \in \mathcal{S}(\mathcal{A}_1^N)$ is the tensor product of N copies of ω , defined by (linear and continuous) extension of

$$(11) \quad \omega^N(B_1 \otimes \dots \otimes B_N) = \omega(B_1) \dots \omega(B_N).$$

This limit exists by definition of an approximately symmetric sequence.²⁶²

It is not difficult to prove that this choice of continuous sections indeed defines a continuous field of C^* -algebras over $I = 0 \cup 1/\mathbb{N}$ with fibers (2). The main point is that

$$(12) \quad \lim_{N \rightarrow \infty} \|A_N\| = \|A_0\|$$

whenever (A_0, A_1, A_2, \dots) satisfies the two conditions above.²⁶³ This is easy to show for symmetric sequences,²⁶⁴ and follows from this for approximately symmetric ones.

Consistent with (8), we conclude that in the limit $N \rightarrow \infty$ the macroscopic observables organize themselves in a commutative C^* -algebra isomorphic to $C(\mathcal{S}(\mathcal{A}_1))$.

6.2 Quasilocal observables

In the C^* -algebraic approach to quantum theory, infinite systems are usually described by means of inductive limit C^* -algebras and the associated quasilocal

²⁶¹A symmetric sequence is evidently approximately symmetric.

²⁶²If (A_1, A_2, \dots) is symmetric with (7), one has $\omega^N(A_N) = \omega^M(A_M)$ for $N > M$, so that the tail of the sequence $(\omega^N(A_N))$ is even independent of N . In the approximately symmetric case one easily proves that $(\omega^N(A_N))$ is a Cauchy sequence.

²⁶³Given (12), the claim follows from Prop. II.1.2.3 in Landsman [1998] and the fact that the set of functions A_0 on $\mathcal{S}(\mathcal{A}_1)$ arising in the said way are dense in $C(\mathcal{S}(\mathcal{A}_1))$ (equipped with the supremum-norm). This follows from the Stone-Weierstrass theorem, from which one infers that the functions in question even exhaust $\mathcal{S}(\mathcal{A}_1)$.

²⁶⁴Assume (7), so that $\|A_N\| = \|j_{NN}(A_N)\|$ for $N \geq M$. By the C^* -axiom $\|A^*A\| = \|A^2\|$ it suffices to prove (12) for $A_0^* = A_0$, which implies $A_M^* = A_M$ and hence $A_N^* = A_N$ for all $N \geq M$. One then has $\|A_N\| = \sup\{|\rho(A_N)|, \rho \in \mathcal{S}(\mathcal{A}_1^N)\}$. Because of the special form of A_N one may replace the supremum over the set $\mathcal{S}(\mathcal{A}_1^N)$ of all states on \mathcal{A}_1^N by the supremum over the set $\mathcal{S}^p(\mathcal{A}_1^N)$ of all permutation invariant states, which in turn may be replaced by the supremum over the extreme boundary $\partial\mathcal{S}^p(\mathcal{A}_1^N)$ of $\mathcal{S}^p(\mathcal{A}_1^N)$. It is well known ([Størmer, 1969]; see also Subsection 6.2) that the latter consists of all states of the form $\rho = \omega^N$, so that $\|A_N\| = \sup\{|\omega^N(A_N)|, \omega \in \mathcal{S}(\mathcal{A}_1)\}$. This is actually equal to $\|A_M\| = \sup\{|\omega^M(A_M)|\}$. Now the norm in $\mathcal{A}_0^{(e)}$ is $\|A_0\| = \sup\{|A_0(\omega)|, \omega \in \mathcal{S}(\mathcal{A}_1)\}$, and by definition of A_0 one has $A_0(\omega) = \omega^M(A_M)$. Hence (12) follows.

observables [Thirring, 1983; Strocchi, 1985; Bratteli and Robinson, 1981; 1987; Haag, 1992; Araki, 1999; Sewell, 1986; 2002]. To arrive at these notions in the case at hand, we proceed as follows [Duffield and Werner, 1992c].

A sequence $A = (A_1, A_2, \dots)$ (where $A_N \in \mathcal{A}_1^N$, as before) is called *local* when for some fixed M and all $N \geq M$ one has $A_N = A_M \otimes 1 \otimes \dots \otimes 1$ (where one has $N - M$ copies of the unit $1 \in \mathcal{A}_1$); cf. (4). A sequence is said to be *quasiloca*l when for any $\varepsilon > 0$ there is an N_ε and a local sequence A' such that $\|A_N - A'_N\| < \varepsilon$ for all $N \geq N_\varepsilon$. On this basis, we define the *inductive limit* C^* -algebra

$$(13) \quad \overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$$

of the family of C^* -algebras (\mathcal{A}_1^N) with respect to the inclusion maps $\mathcal{A}_1^N \hookrightarrow \mathcal{A}_1^{N+1}$ given by $A_N \mapsto A_N \otimes 1$. As a set, (13) consists of all equivalence classes $[A] \equiv A_0$ of quasilocal sequences A under the equivalence relation $A \sim B$ when $\lim_{N \rightarrow \infty} \|A_N - B_N\| = 0$. The norm on $\overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$ is

$$(14) \quad \|A_0\| = \lim_{N \rightarrow \infty} \|A_N\|,$$

and the rest of the C^* -algebraic structure is inherited from the quasilocal sequences in the obvious way (e.g., $A_0^* = [A^*]$ with $A^* = (A_1^*, A_2^*, \dots)$, etc.). As the notation suggests, each \mathcal{A}_1^N is contained in $\overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$ as a C^* -subalgebra by identifying $A_N \in \mathcal{A}_1^N$ with the local (and hence quasilocal) sequence $A = (0, \dots, 0, A_N \otimes 1, A_N \otimes 1 \otimes 1, \dots)$, and forming its equivalence class A_0 in $\overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$ as just explained.²⁶⁵ The assumption underlying the common idea that (13) is “the” algebra of observables of the infinite system under study is that by locality or some other human limitation the infinite tail of the system is not accessible, so that the observables must be arbitrarily close (i.e. in norm) to operators of the form $A_N \otimes 1 \otimes 1, \dots$ for some *finite* N .

This leads us to a second continuous field of C^* -algebras $\mathcal{A}^{(q)}$ over $0 \cup 1/\mathbb{N}$, with fibers

$$(15) \quad \begin{aligned} \mathcal{A}_0^{(q)} &= \overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}; \\ \mathcal{A}_{1/N}^{(q)} &= \mathcal{A}_1^N. \end{aligned}$$

Thus the suffix q reminds one of that fact that the limit algebra $\mathcal{A}_0^{(q)}$ consists of quasilocal or quantum-mechanical observables. We equip the collection of C^* -algebras (15) with the structure of a continuous field of C^* -algebras $\mathcal{A}^{(q)}$ over $0 \cup 1/\mathbb{N}$ by declaring that the continuous sections are of the form (A_0, A_1, A_2, \dots) where (A_1, A_2, \dots) is quasilocal and A_0 is defined by this quasilocal sequence as just explained.²⁶⁶ For $N < \infty$ this field has the same fibers

$$(16) \quad \mathcal{A}_{1/N}^{(q)} = \mathcal{A}_{1/N}^{(c)} = \mathcal{A}_1^N$$

²⁶⁵Of course, the entries A_1, \dots, A_{N-1} , which have been put to zero, are arbitrary.

²⁶⁶The fact that this defines a continuous field follows from (14) and Prop. II.1.2.3 in [Landsman, 1998]; cf. footnote 263.

as the continuous field \mathcal{A} of the previous subsection, but the fiber $\mathcal{A}_0^{(q)}$ is completely different from $\mathcal{A}_0^{(c)}$. In particular, if \mathcal{A}_1 is noncommutative then so is $\mathcal{A}_0^{(q)}$, for it contains all \mathcal{A}_1^N .

The relationship between the continuous fields of C^* -algebras $\mathcal{A}^{(q)}$ and $\mathcal{A}^{(c)}$ may be studied in two different (but related) ways. First, we may construct concrete representations of all C^* -algebras \mathcal{A}_1^N , $N < \infty$, as well as of $\mathcal{A}_0^{(c)}$ and $\mathcal{A}_0^{(q)}$ on a single Hilbert space; this approach leads to superselections rules in the traditional sense. This method will be taken up in the next subsection. Second, we may look at those families of states $(\omega_1, \omega_{1/2}, \dots, \omega_{1/N}, \dots)$ (where $\omega_{1/N}$ is a state on \mathcal{A}_1^N) that admit limit states $\omega_0^{(c)}$ and $\omega_0^{(q)}$ on $\mathcal{A}_0^{(c)}$ and $\mathcal{A}_0^{(q)}$, respectively, such that the ensuing families of states $(\omega_0^{(c)}, \omega_1, \omega_{1/2}, \dots)$ and $(\omega_0^{(q)}, \omega_1, \omega_{1/2}, \dots)$ are *continuous* fields of states on $\mathcal{A}^{(c)}$ and on $\mathcal{A}^{(q)}$, respectively (cf. the end of Subsection 5.1).

Now, any state $\omega_0^{(q)}$ on $\mathcal{A}_0^{(q)}$ defines a state $\omega_{0|1/N}^{(q)}$ on \mathcal{A}_1^N by restriction, and the ensuing field of states on $\mathcal{A}^{(q)}$ is clearly continuous. Conversely, any continuous field $(\omega_0^{(q)}, \omega_1, \omega_{1/2}, \dots, \omega_{1/N}, \dots)$ of states on $\mathcal{A}^{(q)}$ becomes arbitrarily close to a field of the above type for N large.²⁶⁷ However, the restrictions $\omega_{0|1/N}^{(q)}$ of a given state $\omega_0^{(q)}$ on $\mathcal{A}_0^{(q)}$ to \mathcal{A}_1^N may not converge to a state $\omega_0^{(c)}$ on $\mathcal{A}_0^{(c)}$ for $N \rightarrow \infty$.²⁶⁸ States $\omega_0^{(q)}$ on $\bigcup_{N \in \mathbb{N}} \mathcal{A}_1^N$ that do have this property will here be called *classical*. In other words, $\omega_{0|1/N}^{(q)}$ is classical when there exists a probability measure μ_0 on $\mathcal{S}(\mathcal{A}_1)$ such that

$$(17) \quad \lim_{N \rightarrow \infty} \int_{\mathcal{S}(\mathcal{A}_1)} d\mu_0(\rho) (\rho^N(A_N) - \omega_{0|1/N}^{(q)}(A_N)) = 0$$

for each (approximately) symmetric sequence (A_1, A_2, \dots) . To analyze this notion we need a brief intermezzo on general C^* -algebras and their representations.

- A *folium* in the state space $\mathcal{S}(\mathcal{B})$ of a C^* -algebra \mathcal{B} is a convex, norm-closed subspace \mathcal{F} of $\mathcal{S}(\mathcal{B})$ with the property that if $\omega \in \mathcal{F}$ and $B \in \mathcal{B}$ such that $\omega(B^*B) > 0$, then the “reduced” state $\omega_B : A \mapsto \omega(B^*AB)/\omega(B^*B)$ must be in \mathcal{F} [Haag *et al.*, 1970].²⁶⁹ For example, if π is a representation of \mathcal{B} on a Hilbert space \mathcal{H} , then the set of all density matrices on \mathcal{H} (i.e. the π -normal states on \mathcal{B})²⁷⁰ comprises a folium \mathcal{F}_π . In particular, each state ω on \mathcal{B} defines a folium $\mathcal{F}_\omega \equiv \mathcal{F}_{\pi_\omega}$ through its GNS-representation π_ω .
- Two representations π and π' are called *disjoint*, written $\pi \perp \pi'$, if no subrepresentation of π is (unitarily) equivalent to a subrepresentation of π' and

²⁶⁷For any fixed quasilocal sequence (A_1, A_2, \dots) and $\varepsilon > 0$, there is an N_ε such that $|\omega_{1/N}(A_N) - \omega_{0|1/N}^{(q)}(A_N)| < \varepsilon$ for all $N > N_\varepsilon$.

²⁶⁸See footnote 288 below for an example

²⁶⁹See also [Haag, 1992]. The name ‘folium’ is very badly chosen, since $\mathcal{S}(\mathcal{B})$ is by no means foliated by its folia; for example, a folium may contain subfolia.

²⁷⁰A state ω on \mathcal{B} is called π -normal when it is of the form $\omega(B) = \text{Tr } \rho \pi(B)$ for some density matrix ρ . Hence the π -normal states are the normal states on the von Neumann algebra $\pi(\mathcal{B})''$.

vice versa. They are said to be *quasi-equivalent*, written $\pi \sim \pi'$, when π has no subrepresentation disjoint from π' , and vice versa.²⁷¹ Quasi-equivalence is an equivalence relation \sim on the set of representations. See [Kadison and Ringrose, 1986, Ch. 10].

- Similarly, two states ρ, σ are called either quasi-equivalent ($\rho \sim \sigma$) or disjoint ($\rho \perp \sigma$) when the corresponding GNS-representations have these properties.
- A state ω is called *primary* when the corresponding von Neumann algebra $\pi_\omega(\mathcal{B})''$ is a factor.²⁷² Equivalently, ω is primary iff each subrepresentation of $\pi_\omega(\mathcal{B})$ is quasi-equivalent to $\pi_\omega(\mathcal{B})$, which is the case iff $\pi_\omega(\mathcal{B})$ admits no (nontrivial) decomposition as the direct sum of two disjoint subrepresentations.

Now, there is a bijective correspondence between folia in $\mathcal{S}(\mathcal{B})$ and quasi-equivalence classes of representations of \mathcal{B} , in that $\mathcal{F}_\pi = \mathcal{F}_{\pi'}$ iff $\pi \sim \pi'$. Furthermore (as one sees from the GNS-construction), any folium $\mathcal{F} \subset \mathcal{S}(\mathcal{B})$ is of the form $\mathcal{F} = \mathcal{F}_\pi$ for some representation $\pi(\mathcal{B})$. Note that if π is injective (i.e. faithful), then the corresponding folium is dense in $\mathcal{S}(\mathcal{B})$ in the weak*-topology by Fell's Theorem. So in case that \mathcal{B} is simple,²⁷³ any folium is weak*-dense in the state space.

Two states need not be either disjoint or quasi-equivalent. This dichotomy does apply, however, within the class of primary states. Hence *two primary states are either disjoint or quasi-equivalent*. If ω is primary, then each state in the folium of π_ω is primary as well, and is quasi-equivalent to ω . If, on the other hand, ρ and σ are primary and disjoint, then $\mathcal{F}_\rho \cap \mathcal{F}_\sigma = \emptyset$. Pure states are, of course, primary.²⁷⁴ Furthermore, in thermodynamics pure phases are described by primary KMS states [Emch and Knops, 1970; Bratteli and Robinson, 1981; Haag, 1992; Sewell, 2002]. This apparent relationship between primary states and “purity” of some sort is confirmed by our description of macroscopic observables:²⁷⁵

- If $\omega_0^{(q)}$ is a classical primary state on $\mathcal{A}_0^{(q)} = \overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$, then the corresponding limit state $\omega_0^{(c)}$ on $\mathcal{A}_0^{(c)} = C(\mathcal{S}(\mathcal{A}_1))$ is pure (and hence given by a point in $\mathcal{S}(\mathcal{A}_1)$).

²⁷¹Equivalently, two representations π and π' are disjoint iff no π -normal state is π' -normal and vice versa, and quasi-equivalent iff each π -normal state is π' -normal and vice versa.

²⁷²A von Neumann algebra \mathcal{M} acting on a Hilbert space is called a *factor* when its center $\mathcal{M} \cap \mathcal{M}'$ is trivial, i.e. consists of multiples of the identity.

²⁷³In the sense that it has no *closed* two-sided ideals. For example, the matrix algebra $M_n(\mathbb{C})$ is simple for any n , as is its infinite-dimensional analogue, the C^* -algebra of all compact operators on a Hilbert space. The C^* -algebra of quasilocal observables of an infinite quantum systems is typically simple as well.

²⁷⁴Since the corresponding GNS-representation π_ω is irreducible, $\pi_\omega(\mathcal{B})'' = \mathcal{B}(\mathcal{H}_\omega)$ is a factor.

²⁷⁵These claims easily follow from Sewell [2002], §2.6.5, which in turn relies on Hepp [1972].

- If $\rho_0^{(q)}$ and $\sigma_0^{(q)}$ are classical primary states on $\mathcal{A}_0^{(q)}$, then

$$(18) \quad \rho_0^{(c)} = \sigma_0^{(c)} \iff \rho_0^{(q)} \sim \sigma_0^{(q)};$$

$$(19) \quad \rho_0^{(c)} \neq \sigma_0^{(c)} \iff \rho_0^{(q)} \perp \sigma_0^{(q)}.$$

As in (17), a general classical state $\omega_0^{(q)}$ with limit state $\omega_0^{(c)}$ on $C(\mathcal{S}(\mathcal{A}_1))$ defines a probability measure μ_0 on $\mathcal{S}(\mathcal{A}_1)$ by

$$(20) \quad \omega_0^{(c)}(f) = \int_{\mathcal{S}(\mathcal{A}_1)} d\mu_0 f,$$

which describes the probability distribution of the macroscopic observables in that state. As we have seen, this distribution is a delta function for primary states. In any case, it is insensitive to the microscopic details of $\omega_0^{(q)}$ in the sense that local modifications of $\omega_0^{(q)}$ do not affect the limit state $\omega_0^{(c)}$ [Sewell, 2002]. Namely, it easily follows from (8) and the fact that the GNS-representation is cyclic that one can strengthen the second claim above:

Each state in the folium $\mathcal{F}_{\omega_0^{(q)}}$ of a classical state $\omega_0^{(q)}$ is automatically classical and has the same limit state on $\mathcal{A}_0^{(c)}$ as $\omega_0^{(q)}$.

To make this discussion a bit more concrete, we now identify an important class of classical states on $\overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$. We say that a state ω on this C^* -algebra is *permutation-invariant* when each of its restrictions to \mathcal{A}_1^N is invariant under the natural action of the symmetric group \mathfrak{S}_N on \mathcal{A}_1^N (i.e. $\sigma \in \mathfrak{S}_N$ maps an elementary tensor $A_N = B_1 \otimes \cdots \otimes B_N \in \mathcal{A}_1^N$ to $B_{\sigma(1)} \otimes \cdots \otimes B_{\sigma(N)}$, cf. (5)). The structure of the set $\mathcal{S}^\mathfrak{S}$ of all permutation-invariant states in $\mathcal{S}(\mathcal{A}_0^{(q)})$ has been analyzed by Størmer [1969]. Like any compact convex set, it is the (weak*-closed) convex hull of its extreme boundary $\partial_e \mathcal{S}^\mathfrak{S}$. The latter consists of all infinite product states $\omega = \rho^\infty$, where $\rho \in \mathcal{S}(\mathcal{A}_1)$. I.e. if $A_0 \in \mathcal{A}_0^{(q)}$ is an equivalence class $[A_1, A_2, \dots]$, then

$$(21) \quad \rho^\infty(A_0) = \lim_{N \rightarrow \infty} \rho^N(A_N);$$

cf. (11). Equivalently, the restriction of ω to any $\mathcal{A}_1^N \subset \mathcal{A}_0^{(q)}$ is given by $\otimes^N \rho$. Hence $\partial_e \mathcal{S}^\mathfrak{S}$ is isomorphic (as a compact convex set) to $\mathcal{S}(\mathcal{A}_1)$ in the obvious way, and the primary states in $\mathcal{S}^\mathfrak{S}$ are precisely the elements of $\partial_e \mathcal{S}^\mathfrak{S}$.

A general state $\omega_0^{(q)}$ in $\mathcal{S}^\mathfrak{S}$ has a unique decomposition²⁷⁶

$$(22) \quad \omega_0^{(q)}(A_0) = \int_{\mathcal{S}(\mathcal{A}_1)} d\mu(\rho) \rho^\infty(A_0),$$

where μ is a probability measure on $\mathcal{S}(\mathcal{A}_1)$ and $A_0 \in \mathcal{A}_0^{(q)}$.²⁷⁷ The following beautiful illustration of the abstract theory [Unnerstall, 1990a,b] is then clear from (17) and (22):

*If $\omega_0^{(q)}$ is permutation-invariant, then it is classical. The associated limit state $\omega_0^{(c)}$ on $\mathcal{A}_0^{(c)}$ is characterized by the fact that the measure μ_0 in (20) coincides with the measure μ in (22).*²⁷⁸

6.3 Superselection rules

Infinite quantum systems are often associated with the notion of a superselection rule (or sector), which was originally introduced by Wick, Wightman and Wigner [1952] in the setting of standard quantum mechanics on a Hilbert space \mathcal{H} . The basic idea may be illustrated in the example of the boson/fermion (or “univalence”) superselection rule.²⁷⁹ Here one has a *projective* unitary representation \mathcal{D} of the rotation group $SO(3)$ on \mathcal{H} , for which $\mathcal{D}(R_{2\pi}) = \pm 1$ for any rotation $R_{2\pi}$ of 2π around some axis. Specifically, on bosonic states Ψ_B one has $\mathcal{D}(R_{2\pi})\Psi_B = \Psi_B$, whereas on fermionic states Ψ_F the rule is $\mathcal{D}(R_{2\pi})\Psi_F = -\Psi_F$. Now the argument is that a rotation of 2π accomplishes nothing, so that it cannot change the physical state of the system. This requirement evidently holds on the subspace $\mathcal{H}_B \subset \mathcal{H}$ of bosonic states in \mathcal{H} , but it is equally well satisfied on the subspace $\mathcal{H}_F \subset \mathcal{H}$ of fermionic states, since Ψ and $z\Psi$ with $|z| = 1$ describe the same physical state. However, if $\Psi = c_B\Psi_B + c_F\Psi_F$ (with $|c_B|^2 + |c_F|^2 = 1$), then $\mathcal{D}(R_{2\pi})\Psi = c_B\Psi_B - c_F\Psi_F$, which is not proportional to Ψ and apparently describes a genuinely different physical state from Ψ .

The way out is to deny this conclusion by declaring that $\mathcal{D}(R_{2\pi})\Psi$ and Ψ *do* describe the same physical state, and this is achieved by postulating that no phys-

²⁷⁶This follows because $\mathcal{S}^\mathfrak{S}$ is a so-called Bauer simplex [Alfsen, 1970]. This is a compact convex set K whose extreme boundary $\partial_e K$ is closed and for which every $\omega \in K$ has a *unique* decomposition as a probability measure supported by $\partial_e K$, in the sense that $a(\omega) = \int_{\partial_e K} d\mu(\rho) a(\rho)$ for any continuous affine function a on K . For a unital C^* -algebra \mathcal{A} the continuous affine functions on the state space $K = \mathcal{S}(\mathcal{A})$ are precisely the elements A of \mathcal{A} , reinterpreted as functions \hat{A} on $\mathcal{S}(\mathcal{A})$ by $\hat{A}(\omega) = \omega(A)$. For example, the state space $\mathcal{S}(\mathcal{A})$ of a commutative unital C^* -algebra \mathcal{A} is a Bauer simplex, which consists of all (regular Borel) probability measures on the pre state space $\mathcal{P}(\mathcal{A})$.

²⁷⁷This is a quantum analogue of De Finetti’s representation theorem in classical probability theory [Heath and Sudderth, 1976; van Fraassen, 1991]; see also [Hudson and Moody, 1975/76] and [Caves *et al.*, 2002].

²⁷⁸In fact, each state in the folium $\mathcal{F}^\mathfrak{S}$ in $\mathcal{S}(\mathcal{A}_0^{(q)})$ corresponding to the (quasi-equivalence class of) the representation $\oplus_{[\omega \in \mathcal{S}^\mathfrak{S}]} \pi_\omega$ is classical.

²⁷⁹See also [Giulini, 2003] for a modern mathematical treatment.

ical *observables* A (in their usual mathematical guise as operators on \mathcal{H}) exist for which $(\Psi_B, A\Psi_F) \neq 0$. For in that case one has

$$(23) \quad (c_B\Psi_B \pm c_F\Psi_F, A(c_B\Psi_B \pm c_F\Psi_F)) = |c_B|^2(\Psi_B, A\Psi_B) + |c_F|^2(\Psi_F, A\Psi_F)$$

for any *observable* A , so that $(\mathcal{D}(R_{2\pi})\Psi, A\mathcal{D}(R_{2\pi})\Psi) = (\Psi, A\Psi)$ for any $\Psi \in \mathcal{H}$. Since any quantum-mechanical prediction ultimately rests on expectation values $(\Psi, A\Psi)$ for physical observables A , the conclusion is that a rotation of 2π indeed does nothing to the system. This is codified by saying that superpositions of the type $c_B\Psi_B + c_F\Psi_F$ are *incoherent* (whereas superpositions $c_1\Psi_1 + c_2\Psi_2$ with Ψ_1, Ψ_2 both in either \mathcal{H}_B or in \mathcal{H}_F are *coherent*). Each of the subspaces \mathcal{H}_B and \mathcal{H}_F of \mathcal{H} is said to be a *superselection sector*, and the statement that $(\Psi_B, A\Psi_F) = 0$ for any observable A and $\Psi_B \in \mathcal{H}_B$ and $\Psi_F \in \mathcal{H}_F$ is called a *superselection rule*.²⁸⁰

The price one pays for this solution is that states of the form $c_B\Psi_B + c_F\Psi_F$ with $c_B \neq 0$ and $c_F \neq 0$ are mixed, as one sees from (23). More generally, if $\mathcal{H} = \oplus_{\lambda \in \Lambda} \mathcal{H}_\lambda$ with $(\Psi, A\Phi) = 0$ whenever A is an observable, $\Psi \in \mathcal{H}_\lambda$, $\Phi \in \mathcal{H}_{\lambda'}$, and $\lambda \neq \lambda'$, and if in addition for each λ and each pair $\Psi, \Phi \in \mathcal{H}_\lambda$ there exists an observable A for which $(\Psi, A\Phi) \neq 0$, then the subspaces \mathcal{H}_λ are called superselection sectors in \mathcal{H} . Again a key consequence of the occurrence of superselection sectors is that unit vectors of the type $\Psi = \sum_\lambda c_\lambda \Psi_\lambda$ with $\Psi \in \mathcal{H}_\lambda$ (and $c_\lambda \neq 0$ for at least two λ 's) define mixed states

$$\psi(A) = (\Psi, A\Psi) = \sum_\lambda |c_\lambda|^2 (\Psi_\lambda, A\Psi_\lambda) = \sum_\lambda |c_\lambda|^2 \psi_\lambda(A).$$

This procedure is rather ad hoc. A much deeper approach to superselection theory was developed by Haag and collaborators; see [Roberts and Roepstorff, 1969] for an introduction. Here the starting point is the abstract C^* -algebra of observables \mathcal{A} of a given quantum system, and superselection sectors are reinterpreted as equivalence classes (under unitary isomorphism) of irreducible representations of \mathcal{A} (satisfying a certain selection criterion — see below). The connection between the concrete Hilbert space approach to superselection sectors discussed above and the abstract C^* -algebraic approach is given by the following lemma [Hepp, 1972]:²⁸¹

Two pure states ρ, σ on a C^ -algebra \mathcal{A} define different sectors iff for each representation $\pi(\mathcal{A})$ on a Hilbert space \mathcal{H} containing unit vectors Ψ_ρ, Ψ_σ such that $\rho(A) = (\Psi_\rho, \pi(A)\Psi_\rho)$ and $\sigma(A) = (\Psi_\sigma, \pi(A)\Psi_\sigma)$ for all $A \in \mathcal{A}$, one has $(\Psi_\rho, \pi(A)\Psi_\sigma) = 0$ for all $A \in \mathcal{A}$.*

In practice, however, most irreducible representations of a typical C^* -algebra \mathcal{A} used in physics are physically irrelevant mathematical artefacts. Such representations may be excluded from consideration by some *selection criterion*. What this

²⁸⁰In an ordinary selection rule between Ψ and Φ one merely has $(\Psi, H\Phi) = 0$ for the Hamiltonian H .

²⁸¹Hepp proved a more general version of this lemma, in which ‘Two pure states ρ, σ on a C^* -algebra \mathcal{B} define different sectors iff. . .’ is replaced by ‘Two states ρ, σ on a C^* -algebra \mathcal{B} are disjoint iff. . .’

means depends on the context. For example, in quantum field theory this notion is made precise in the so-called DHR theory (reviewed by Roberts [1990], Haag [1992], Araki [1999], and Halvorson [2005]). In the class of theories discussed in the preceding two subsections, we take the algebra of observables \mathcal{A} to be $\mathcal{A}_0^{(q)}$ — essentially for reasons of human limitation — and for pedagogical reasons define (equivalence classes of) irreducible representations of $\mathcal{A}_0^{(q)}$ as superselection sectors, henceforth often just called *sectors*, only when they are equivalent to the GNS-representation given by a permutation-invariant pure state on $\mathcal{A}_0^{(q)}$. In particular, such a state is classical. On this selection criterion, the results in the preceding subsection trivially imply that there is a bijective correspondence between pure states on \mathcal{A}_1 and sectors of $\mathcal{A}_0^{(q)}$. The sectors of the commutative C^* -algebra $\mathcal{A}_0^{(c)}$ are just the points of $\mathcal{S}(\mathcal{A}_1)$; note that a *mixed* state on \mathcal{A}_1 defines a *pure* state on $\mathcal{A}_0^{(c)}$! The role of the sectors of \mathcal{A}_1 in connection with those of $\mathcal{A}_0^{(c)}$ will be clarified in Subsection 6.5.

Whatever the model or the selection criterion, it is enlightening (and to some extent even in accordance with experimental practice) to consider superselection sectors entirely from the perspective of the pure states on the algebra of observables \mathcal{A} , removing \mathcal{A} itself and its representations from the scene. To do so, we equip the space $\mathcal{P}(\mathcal{A})$ of pure states on \mathcal{A} with the structure of a transition probability space [von Neumann, 1981; Mielnik, 1968].²⁸² A *transition probability* on a set \mathcal{P} is a function

$$(24) \quad p : \mathcal{P} \times \mathcal{P} \rightarrow [0, 1]$$

that satisfies

$$(25) \quad p(\rho, \sigma) = 1 \iff \rho = \sigma$$

and

$$(26) \quad p(\rho, \sigma) = 0 \iff p(\sigma, \rho) = 0.$$

A set with such a transition probability is called a *transition probability space*. Now, the pure state space $\mathcal{P}(\mathcal{A})$ of a C^* -algebra \mathcal{A} carries precisely this structure if we define²⁸³

$$(27) \quad p(\rho, \sigma) := \inf\{\rho(A) \mid A \in \mathcal{A}, 0 \leq A \leq 1, \sigma(A) = 1\}.$$

To give a more palatable formula, note that since pure states are primary, two pure states ρ, σ are either disjoint ($\rho \perp \sigma$) or else (quasi, hence unitarily) equivalent ($\rho \sim \sigma$). In the first case, (27) yields

$$(28) \quad p(\rho, \sigma) = 0 \quad (\rho \perp \sigma).$$

²⁸²See also [Beltrametti and Cassinelli, 1984] or [Landsman, 1998] for concise reviews.

²⁸³This definition applies to the case that \mathcal{A} is unital; see [Landsman, 1998] for the general case. An analogous formula defines a transition probability on the extreme boundary of any compact convex set.

In the second case it follows from Kadison’s transitivity theorem (cf. Thm. 10.2.6 in [Kadison and Ringrose, 1986]) that the Hilbert space \mathcal{H}_ρ from the GNS-representation $\pi_\rho(\mathcal{A})$ defined by ρ contains a unit vector Ω_σ (unique up to a phase) such that

$$(29) \quad \sigma(A) = (\Omega_\sigma, \pi_\rho(A)\Omega_\sigma).$$

Eq. (27) then leads to the well-known expression

$$(30) \quad p(\rho, \sigma) = |(\Omega_\rho, \Omega_\sigma)|^2 \quad (\rho \sim \sigma).$$

In particular, if \mathcal{A} is commutative, then

$$(31) \quad p(\rho, \sigma) = \delta_{\rho\sigma}.$$

For $\mathcal{A} = M_2(\mathbb{C})$ one obtains

$$(32) \quad p(\rho, \sigma) = \frac{1}{2}(1 + \cos \theta_{\rho\sigma}),$$

where $\theta_{\rho\sigma}$ is the angular distance between ρ and σ (seen as points on the two-sphere $S^2 = \partial_e B^3$, cf. (3) etc.), measured along a great circle.

Superselection sectors may now be defined for any transition probability spaces \mathcal{P} . A family of subsets of \mathcal{P} is called *orthogonal* if $p(\rho, \sigma) = 0$ whenever ρ and σ do not lie in the same subset. The space \mathcal{P} is called *reducible* if it is the union of two (nonempty) orthogonal subsets; if not, it is said to be *irreducible*. A *component* of \mathcal{P} is a subset $\mathcal{C} \subset \mathcal{P}$ such that \mathcal{C} and $\mathcal{P} \setminus \mathcal{C}$ are orthogonal. An irreducible component of \mathcal{P} is called a (*superselection*) *sector*. Thus \mathcal{P} is the disjoint union of its sectors. For $\mathcal{P} = \mathcal{P}(\mathcal{A})$ this reproduces the algebraic definition of a superselection sector (modulo the selection criterion) via the correspondence between states and representations given by the GNS-constructions. For example, in the commutative case $\mathcal{A} \cong C(X)$ each point in $X \cong \mathcal{P}(\mathcal{A})$ is its own little sector.

6.4 A simple example: the infinite spin chain

Let us illustrate the occurrence of superselection sectors in a simple example, where the algebra of observables is $\mathcal{A}_0^{(q)}$ with $\mathcal{A}_1 = M_2(\mathbb{C})$. Let $\mathcal{H}_1 = \mathbb{C}^2$, so that $\mathcal{H}_1^N = \otimes^N \mathbb{C}^2$ is the tensor product of N copies of \mathbb{C}^2 . It is clear that \mathcal{A}_1^N acts on \mathcal{H}_1^N in a natural way (i.e. componentwise). This defines an irreducible representation π_N of \mathcal{A}_1^N , which is indeed its unique irreducible representation (up to unitary equivalence). In particular, for $N < \infty$ the quantum system whose algebra of observables is \mathcal{A}_1^N (such as a chain with N two-level systems) has no superselection rules. We define the $N \rightarrow \infty$ limit “ $(M_2(\mathbb{C}))^\infty$ ” of the C^* -algebras $(M_2(\mathbb{C}))^N$ as the inductive limit $\mathcal{A}_0^{(q)}$ for $\mathcal{A}_1 = M_2(\mathbb{C})$, as introduced in Subsection 6.2; see (13). The definition of “ $\otimes^\infty \mathbb{C}^2$ ” is slightly more involved, as follows [von Neumann, 1938].

For any Hilbert space \mathcal{H}_1 , let Ψ be a sequence (Ψ_1, Ψ_2, \dots) with $\Psi_n \in \mathcal{H}_1$. The space \mathbf{H}_1 of such sequences is a vector space in the obvious way. Now let Ψ and Φ be two such sequences, and write $(\Psi_n, \Phi_n) = \exp(i\alpha_n)|(\Psi_n, \Phi_n)|$. If $\sum_n |\alpha_n| = \infty$,

we define the (pre-) inner product (Ψ, Φ) to be zero. If $\sum_n |\alpha_n| < \infty$, we put $(\Psi, \Phi) = \prod_n (\Psi_n, \Phi_n)$ (which, of course, may still be zero!). The (vector space) quotient of \mathcal{H}_1 by the space of sequences Ψ for which $(\Psi, \Psi) = 0$ can be completed to a Hilbert space \mathcal{H}_1^∞ in the induced inner product, called the *complete* infinite tensor product of the Hilbert space \mathcal{H}_1 (over the index set \mathbb{N}).²⁸⁴ We apply this construction with $\mathcal{H}_1 = \mathbb{C}^2$. If (e_i) is some basis of \mathbb{C}^2 , an orthonormal basis of \mathcal{H}_1^∞ then consists of all different infinite strings $e_{i_1} \otimes \cdots e_{i_n} \otimes \cdots$, where e_{i_n} is e_i regarded as a vector in \mathbb{C}^2 .²⁸⁵ We denote the multi-index (i_1, \dots, i_n, \dots) simply by I , and the corresponding basis vector by e_I .

This Hilbert space \mathcal{H}_1^∞ carries a natural faithful representation π of $\mathcal{A}_0^{(q)}$: if $A_0 \in \mathcal{A}_0^{(q)}$ is an equivalence class $[A_1, A_2, \dots]$, then $\pi(A_0)e_I = \lim_{N \rightarrow \infty} A_N e_i$, where A_N acts on the first N components of e_I and leaves the remainder unchanged.²⁸⁶ Now the point is that although each \mathcal{A}_1^N acts irreducibly on \mathcal{H}_1^N , the representation $\pi(\mathcal{A}_0^{(q)})$ on \mathcal{H}_1^∞ thus constructed is highly reducible. The reason for this is that by definition (quasi-) local elements of $\mathcal{A}_0^{(q)}$ leave the infinite tail of a vector in \mathcal{H}_1^∞ (almost) unaffected, so that vectors with different tails lie in different superselection sectors. Without the quasi-locality condition on the elements of $\mathcal{A}_0^{(q)}$, no superselection rules would arise. For example, in terms of the usual basis

$$(33) \quad \left\{ \uparrow = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \downarrow = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$$

of \mathbb{C}^2 , the vectors $\Psi_\uparrow = \uparrow \otimes \uparrow \cdots \uparrow \cdots$ (i.e. an infinite product of ‘up’ vectors) and $\Psi_\downarrow = \downarrow \otimes \downarrow \cdots \downarrow \cdots$ (i.e. an infinite product of ‘down’ vectors) lie in different sectors. The reason why the inner product $(\Psi_\uparrow, \pi(A)\Psi_\downarrow)$ vanishes for any $A \in \mathcal{A}_0^{(q)}$ is that for local observables A one has $\pi(A) = A_M \otimes 1 \otimes \cdots 1 \cdots$ for some $A_M \in \mathcal{B}(\mathcal{H}_M)$; the inner product in question therefore involves infinitely many factors $(\uparrow, 1 \downarrow) = (\uparrow, \downarrow) = 0$. For quasilocal A the operator $\pi(A)$ might have a small nontrivial tail, but the inner product vanishes nonetheless by an approximation argument.

More generally, elementary analysis shows that $(\Psi_u, \pi(A)\Psi_v) = 0$ whenever $\Psi_u = \otimes^\infty u$ and $\Psi_v = \otimes^\infty v$ for unit vectors $u, v \in \mathbb{C}^2$ with $u \neq v$. The corresponding vector states ψ_u and ψ_v on $\mathcal{A}_0^{(q)}$ (i.e. $\psi_u(A) = (\Psi_u, \pi(A)\Psi_u)$ etc.) are obviously permutation-invariant and hence classical. Identifying $\mathcal{S}(M_2(\mathbb{C}))$ with B^3 , as in (3), the corresponding limit state $(\psi_u)_0$ on $\mathcal{A}_0^{(c)}$ defined by ψ_u is given by (evaluation at) the point $\tilde{u} = (x, y, z)$ of $\partial_e B^3 = S^2$ (i.e. the two-sphere) for which the corresponding density matrix $\rho(\tilde{u})$ is the projection operator onto u . It

²⁸⁴Each fixed $\Psi \in \mathcal{H}_1$ defines an *incomplete* tensor product \mathcal{H}_Ψ^∞ , defined as the closed subspace of \mathcal{H}_1^∞ consisting of all Φ for which $\sum_n |(\Psi_n, \Phi_n) - 1| < \infty$. If \mathcal{H}_1 is separable, then so is \mathcal{H}_Ψ^∞ (in contrast to \mathcal{H}_1^∞ , which is an uncountable direct sum of the \mathcal{H}_Ψ^∞).

²⁸⁵The cardinality of the set of all such strings equals that of \mathbb{R} , so that \mathcal{H}_1^∞ is non-separable, as claimed.

²⁸⁶Indeed, this yields an alternative way of defining $\overline{\cup_{N \in \mathbb{N}} \mathcal{A}_1^N}$ as the norm closure of the union of all \mathcal{A}_1^N acting on \mathcal{H}_1^∞ in the stated way.

follows that ψ_u and ψ_v are disjoint; cf. (19). We conclude that each unit vector $u \in \mathbb{C}^2$ determines a superselection sector π_u , namely the GNS-representation of the corresponding state ψ_u , and that each such sector is realized as a subspace \mathcal{H}_u of \mathcal{H}_1^∞ (viz. $\mathcal{H}_u = \pi(\mathcal{A}_0^{(q)}\Psi_u)$). Moreover, since a permutation-invariant state on $\mathcal{A}_0^{(q)}$ is pure iff it is of the form ψ_u , we have found all superselection sectors of our system. Thus in what follows we may concentrate our attention on the subspace (of \mathcal{H}_1^∞) and subrepresentation (of π)

$$(34) \quad \begin{aligned} \mathcal{H}_\mathfrak{S} &= \bigoplus_{\tilde{u} \in S^2} \mathcal{H}_u; \\ \pi_\mathfrak{S}(\mathcal{A}_0^{(q)}) &= \bigoplus_{\tilde{u} \in S^2} \pi_u(\mathcal{A}_0^{(q)}), \end{aligned}$$

where π_u is simply the restriction of π to $\mathcal{H}_u \subset \mathcal{H}_1^\infty$.

In the presence of superselection sectors one may construct operators that distinguish different sectors whilst being a multiple of the unit in each sector. In quantum field theory these are typically global charges, and in our example the macroscopic observables play this role. To see this, we return to Subsection 6.1. It is not difficult to show that for any approximately symmetric sequence (A_1, A_2, \dots) the limit

$$(35) \quad \bar{A} = \lim_{N \rightarrow \infty} \pi_\mathfrak{S}(A_N)$$

exists in the strong operator topology on $\mathcal{B}(\mathcal{H}_\mathfrak{S})$ [Bona, 1988]. Moreover, if $A_0 \in \mathcal{A}_0^{(c)} = C(\mathcal{S}(\mathcal{A}_1))$ is the function defined by the given sequence,²⁸⁷ then the map $A_0 \mapsto \bar{A}$ defines a faithful representation of $\mathcal{A}_0^{(c)}$ on $\mathcal{H}_\mathfrak{S}$, which we call $\pi_\mathfrak{S}$ as well (by abuse of notation). An easy calculation in fact shows that $\pi_\mathfrak{S}(A_0)\Psi = A_0(\tilde{u})\Psi$ for $\Psi \in \mathcal{H}_u$, or, in other words,

$$(36) \quad \pi_\mathfrak{S}(A_0) = \bigoplus_{\tilde{u} \in S^2} A_0(\tilde{u})1_{\mathcal{H}_u}.$$

Thus the $\pi_\mathfrak{S}(A_0)$ indeed serve as the operators in question.

To illustrate how delicate all this is, it may be interesting to note that even for symmetric sequences the limit $\lim_{N \rightarrow \infty} \pi(A_N)$ does not exist on \mathcal{H}_1^∞ , not even in the strong topology.²⁸⁸ On the positive side, it can be shown that $\lim_{N \rightarrow \infty} \pi(A_N)\Psi$ exists as an element of the von Neumann algebra $\pi(\mathcal{A}_0^{(q)})''$ whenever the vector state ψ defined by Ψ lies in the folium $\mathcal{F}^\mathfrak{S}$ generated by all permutation-invariant states [Bona, 1988; Unnerstall, 1990a].

This observation is part of a general theory of macroscopic observables in the setting of von Neumann algebras [Primas, 1983; Rieckers, 1984; Amann, 1986; 1987; Morchio and Strocchi, 1987; Bona, 1988; 1989; Unnerstall, 1990a; 1990b; Breuer,

²⁸⁷ Recall that $A_0(\omega) = \lim_{N \rightarrow \infty} \omega^N(A_N)$.

²⁸⁸ For example, let us take the sequence $A_N = j_{N1}(\text{diag}(1, -1))$ and the vector $\Psi = \uparrow\downarrow\uparrow\uparrow\uparrow\uparrow\downarrow\downarrow\downarrow\downarrow\downarrow\downarrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow\uparrow \dots$, where a sequence of 2^N factors of \uparrow is followed by 2^{N+1} factors of \downarrow , etc. Then the sequence $\{\pi(A_N)\Psi\}_{N \in \mathbb{N}}$ in \mathcal{H}_1^∞ diverges: the subsequence where N runs over all numbers 2^n with n odd converges to $\frac{1}{3}\Psi$, whereas the subsequence where N runs over all 2^n with n even converges to $-\frac{1}{3}\Psi$.

1994; Atmanspacher *et al.*, 1999], which complements the purely C^* -algebraic approach of Raggio and Werner [1989; 1991], Duffield and Werner [1992a,b,c], and Duffield *et al.* [1992] explained above.²⁸⁹ In our opinion, the latter has the advantage that conceptually the passage to the limit $N \rightarrow \infty$ (and thereby the idealization of a large system as an infinite one) is very satisfactory, especially in our reformulation in terms of continuous fields of C^* -algebras. Here the commutative C^* -algebra $\mathcal{A}_0^{(c)}$ of macroscopic observables of the infinite system is glued to the noncommutative algebras \mathcal{A}_1^N of the corresponding finite systems in a continuous way, and the continuous sections of the ensuing continuous field of C^* -algebras $\mathcal{A}^{(c)}$ exactly describe how *macroscopic* quantum observables of the finite systems converge to classical ones. *Microscopic* quantum observables of the pertinent finite systems, on the other hand, converge to quantum observables of the infinite quantum system, and this convergence is described by the continuous sections of the continuous field of C^* -algebras $\mathcal{A}^{(q)}$. This entirely avoids the language of superselection rules, which rather displays a shocking *discontinuity* between finite and infinite systems: for superselection rules do not exist in finite systems!²⁹⁰

6.5 Poisson structure and dynamics

We now pass to the discussion of time-evolution in infinite systems of the type considered so far. We start with the observation that the state space $\mathcal{S}(\mathcal{B})$ of a finite-dimensional C^* -algebra \mathcal{B} (for simplicity assumed unital in what follows) is a Poisson manifold (cf. Subsection 4.3) in a natural way. A similar statement holds in the infinite-dimensional case, and we carry the reader through the necessary adaptations of the main argument by means of footnotes.²⁹¹ We write $K = \mathcal{S}(\mathcal{B})$.

²⁸⁹Realistic models have been studied in the context of both the C^* -algebraic and the von Neumann algebraic approach by Rieckers and his associates. See, for example, [Honegger and Rieckers, 1994; Gerisch *et al.*, 1999; Gerisch, *et al.* 2003], and many other papers. For altogether different approaches to macroscopic observables see [van Kampen, 1954; 1988; 1993; Wan and Fountain, 1998; Harrison and Wan, 1997; Wan *et al.*, 1998; Fröhlich *et al.*, 2002], and [Poulin, 2004].

²⁹⁰We here refer to superselection rules in the traditional sense of inequivalent irreducible representations of *simple* C^* -algebras. For topological reasons certain finite-dimensional systems are described by (non-simple) C^* -algebras that do admit inequivalent irreducible representations [Landsman, 1990a,b].

²⁹¹Of which this is the first. When \mathcal{B} is infinite-dimensional, the state space $\mathcal{S}(\mathcal{B})$ is no longer a manifold, let alone a Poisson manifold, but a *Poisson space* [Landsman, 1997; 1998]. This is a generalization of a Poisson manifold, which turns a crucial *property* of the latter into a *definition*. This property is the foliation of a Poisson manifold by its symplectic leaves [Weinstein, 1983], and the corresponding definition is as follows: A *Poisson space* P is a Hausdorff space of the form $P = \cup_{\alpha} S_{\alpha}$ (*disjoint union*), where each S_{α} is a symplectic manifold (possibly infinite-dimensional) and each injection $\iota_{\alpha} : S_{\alpha} \hookrightarrow P$ is continuous. Furthermore, one has a linear subspace $F \subset C(P, \mathbb{R})$ that separates points and has the property that the restriction of each $f \in F$ to each S_{α} is smooth. Finally, if $f, g \in F$ then $\{f, g\} \in F$, where the Poisson bracket is defined by $\{f, g\}(\iota_{\alpha}(\sigma)) = \{\iota_{\alpha}^* f, \iota_{\alpha}^* g\}_{\alpha}(\sigma)$. Clearly, a Poisson manifold M defines a Poisson space if one takes $P = M$, $F = C^{\infty}(M)$, and the S_{α} to be the symplectic leaves defined by the given Poisson bracket. Thus we refer to the manifolds S_{α} in the above definition as the *symplectic leaves* of P as well.

Firstly, an element $A \in \mathcal{B}$ defines a linear function \hat{A} on \mathcal{B}^* and hence on K (namely by restriction) through $\hat{A}(\omega) = \omega(A)$. For such functions we define the Poisson bracket by

$$(37) \quad \{\hat{A}, \hat{B}\} = i\widehat{[A, B]}.$$

Here the factor i has been inserted in order to make the Poisson bracket of two real-valued functions real-valued again; for \hat{A} is real-valued on K precisely when A is self-adjoint, and if $A^* = A$ and $B^* = B$, then $i[A, B]$ is self-adjoint (whereas $[A, B]$ is skew-adjoint). In general, for $f, g \in C^\infty(K)$ we put

$$(38) \quad \{f, g\}(\omega) = i\omega([df_\omega, dg_\omega]),$$

interpreted as follows.²⁹² Let $\mathcal{B}_\mathbb{R}$ be the self-adjoint part of \mathcal{B} , and interpret K as a subspace of $\mathcal{B}_\mathbb{R}^*$; since a state ω satisfies $\omega(A^*) = \overline{\omega(A)}$ for all $A \in \mathcal{B}$, it is determined by its values on self-adjoint elements. Subsequently, we identify the tangent space at ω with

$$(39) \quad T_\omega K = \{\rho \in \mathcal{B}_\mathbb{R}^* \mid \rho(1) = 0\} \subset \mathcal{B}_\mathbb{R}^*$$

and the cotangent space at ω with the quotient (of real Banach spaces)

$$(40) \quad T_\omega^* K = \mathcal{B}_\mathbb{R}^{**} / \mathbb{R}1,$$

where the unit $1 \in \mathcal{B}$ is regarded as an element of \mathcal{B}^{**} through the canonical embedding $\mathcal{B} \subset \mathcal{B}^{**}$. Consequently, the differential forms df and dg at $\omega \in K$ define elements of $\mathcal{B}_\mathbb{R}^{**} / \mathbb{R}1$. The commutator in (38) is then defined as follows: one lifts $df_\omega \in \mathcal{B}_\mathbb{R}^{**} / \mathbb{R}1$ to $\mathcal{B}_\mathbb{R}^{**}$, and uses the natural isomorphism $\mathcal{B}^{**} \cong \mathcal{B}$ typical of finite-dimensional vector spaces.²⁹³ The arbitrariness in this lift is a multiple of 1, which drops out of the commutator. Hence $i[df_\omega, dg_\omega]$ is an element of $\mathcal{B}_\mathbb{R}^{**} \cong \mathcal{B}_\mathbb{R}$, on which the value of the functional ω is defined.²⁹⁴ This completes the definition of the Poisson bracket; one easily recovers (37) as a special case of (38).

The symplectic leaves of the given Poisson structure on K have been determined by Duffield and Werner [1992a].²⁹⁵ Namely:

Two states ρ and σ lie in the same symplectic leaf of $\mathcal{S}(\mathcal{B})$ iff $\rho(A) = \sigma(UAU^)$ for some unitary $U \in \mathcal{B}$.*

²⁹²In the infinite-dimensional case $C^\infty(K)$ is defined as the intersection of the smooth functions on K with respect to its Banach manifold structure and the space $C(K)$ of weak*-continuous functions on K . The differential forms df and dg in (38) also require an appropriate definition; see [Duffield and Werner, 1992a; Bona, 2000], and [Odzijewicz and Ratiu, 2003] for the technicalities.

²⁹³In the infinite-dimensional case one uses the canonical identification between \mathcal{B}^{**} and the enveloping von Neumann algebra of \mathcal{B} to define the commutator.

²⁹⁴If \mathcal{B} is infinite-dimensional, one here regards \mathcal{B}^* as the predual of the von Neumann algebra \mathcal{B}^{**} .

²⁹⁵See also [Bona, 2000] for the infinite-dimensional special case where \mathcal{B} is the C^* -algebra of compact operators.

When ρ and σ are pure, this is the case iff the corresponding GNS-representations $\pi_\rho(\mathcal{B})$ and $\pi_\sigma(\mathcal{B})$ are unitarily equivalent,²⁹⁶ but in general the implication holds only in one direction: if ρ and σ lie in the same leaf, then they have unitarily equivalent GNS-representations.²⁹⁷

It follows from this characterization of the symplectic leaves of $K = \mathcal{S}(\mathcal{B})$ that the pure state space $\partial_e K = \mathcal{P}(\mathcal{B})$ inherits the Poisson bracket from K , and thereby becomes a Poisson manifold in its own right.²⁹⁸ This leads to an important connection between the superselection sectors of \mathcal{B} and the Poisson structure on $\mathcal{P}(\mathcal{B})$ [Landsman, 1997; 1998]:

The sectors of the pure state space $\mathcal{P}(\mathcal{B})$ of a C^ -algebra \mathcal{B} as a transition probability space coincide with its symplectic leaves as a Poisson manifold.*

For example, when $\mathcal{B} \cong C(X)$ is commutative, the space $\mathcal{S}(C(X))$ of all (regular Borel) probability measures on X acquires a Poisson bracket that is identically zero, as does its extreme boundary X . It follows from (31) that the sectors in X are its points, and so are its symplectic leaves (in view of their definition and the vanishing Poisson bracket). The simplest noncommutative case is $\mathcal{B} = M_2(\mathbb{C})$, for which the symplectic leaves of the state space $K = \mathcal{S}(M_2(\mathbb{C})) \cong B^3$ (cf. (3)) are the spheres with constant radius.²⁹⁹ The sphere with radius 1 consists of points in B^3 that correspond to pure states on $M_2(\mathbb{C})$, all interior symplectic leaves of K coming from mixed states on $M_2(\mathbb{C})$.

The coincidence of sectors and symplectic leaves of $\mathcal{P}(\mathcal{B})$ is a compatibility condition between the transition probability structure and the Poisson structure. It is typical of the specific choices (27) and (38), respectively, and hence of quantum theory. In classical mechanics one has the freedom of equipping a manifold M with an arbitrary Poisson structure, and yet use $C_0(M)$ as the commutative C^* -algebra of observables. The transition probability (31) (which follows from (27) in the commutative case) are clearly the correct ones in classical physics, but since the

²⁹⁶Cf. Thm. 10.2.6 in [Kadison and Ringrose, 1986].

²⁹⁷An important step of the proof is the observation that the Hamiltonian vector field $\xi_f(\omega) \in T_\omega K \subset \mathcal{A}_\mathbb{R}^*$ of $f \in C^\infty(K)$ is given by $\langle \xi_f(\omega), B \rangle = i[df_\omega, B]$, where $B \in \mathcal{B}_\mathbb{R} \subset \mathcal{B}_\mathbb{R}^{**}$ and $df_\omega \in \mathcal{B}_\mathbb{R}^{**}/\mathbb{R}1$. (For example, this gives $\xi_{\hat{A}} \hat{B} = i[\widehat{A}, \widehat{B}] = \{\hat{A}, \hat{B}\}$ by (37), as it should be.) If φ_t^h denotes the Hamiltonian flow of h at time t , it follows (cf. [Duffield *et al.*, 1992, Prop. 6.1] or [Duffield and Werner, 1992a, Prop. 3.1] that $\langle \varphi_t^h(\omega), B \rangle = \langle \omega, U_t^h B (U_t^h)^* \rangle$ for some unitary $U_t^h \in \mathcal{B}$. For example, if $h = \hat{A}$ then $U_t^h = \exp(itA)$.

²⁹⁸More generally, a Poisson space. The structure of $\mathcal{P}(\mathcal{B})$ as a Poisson space was introduced by Landsman [1997; 1998] without recourse to the full state space or the work of Duffield and Werner [1992a].

²⁹⁹Equipped with a multiple of the so-called Fubini–Study symplectic structure; see [Landsman, 1998] or any decent book on differential geometry for this notion. This claim is immediate from footnote 260. More generally, the pure state space of $M_n(\mathbb{C})$ is the projective space $\mathbb{P}\mathbb{C}^n$, which again becomes equipped with the Fubini–Study symplectic structure. This is even true for $n = \infty$ if one defines $M_\infty(\mathbb{C})$ as the C^* -algebra of compact operators on a separable Hilbert space \mathcal{H} : in that case one has $\mathcal{P}(M_\infty(\mathbb{C})) \cong \mathbb{P}\mathcal{H}$. Cf. [Cantoni, 1977; Cirelli *et al.*, 1983; Cirelli *et al.*, 1990; Landsman, 1998; Ashtekar and Schilling, 1999; Marmo *et al.*, 2005], etc.

symplectic leaves of M can be almost anything, the coincidence in question does not hold.

However, there exists a compatibility condition between the transition probability structure and the Poisson structure, which is shared by classical and quantum theory. This is the property of *unitarity* of a Hamiltonian flow, which in the present setting we formulate as follows.³⁰⁰ First, in quantum theory with algebra of observables \mathcal{B} we define time-evolution (in the sense of an automorphic action of the abelian group \mathbb{R} on \mathcal{B} , i.e. a one-parameter group α of automorphisms on \mathcal{B}) to be *Hamiltonian* when $A(t) = \alpha_t(A)$ satisfies the Heisenberg equation $i\hbar dA/dt = [A, H]$ for some self-adjoint element $H \in \mathcal{B}$. The corresponding flow on $\mathcal{P}(\mathcal{B})$ — i.e. $\omega_t(A) = \omega(A(t))$ — is equally well said to be Hamiltonian in that case. In classical mechanics with Poisson manifold M we similarly say that a flow on M is Hamiltonian when it is the flow of a Hamiltonian vector field ξ_h for some $h \in C^\infty(M)$. (Equivalently, the time-evolution of the observables $f \in C^\infty(M)$ is given by $df/dt = \{h, f\}$; cf. (18) etc.) The point is that in either case the flow is unitary in the sense that

$$(41) \quad p(\rho(t), \sigma(t)) = p(\rho, \sigma)$$

for all t and all $\rho, \sigma \in P$ with $P = \mathcal{P}(\mathcal{B})$ (equipped with the transition probabilities (27) and the Poisson bracket (38)) or $P = M$ (equipped with the transition probabilities (31) and any Poisson bracket).³⁰¹

In both cases $P = \mathcal{P}(\mathcal{B})$ and $P = M$, a Hamiltonian flow has the property (which is immediate from the definition of a symplectic leaf) that for all (finite) times t a point $\omega(t)$ lies in the same symplectic leaf of P as $\omega = \omega(0)$. In particular, in quantum theory $\omega(t)$ and ω must lie in the same sector. In the quantum theory of infinite systems an automorphic time-evolution is rarely Hamiltonian, but one reaches a similar conclusion under a weaker assumption. Namely, if a given one-parameter group of automorphisms α on \mathcal{B} is *implemented* in the GNS-representation $\pi_\omega(\mathcal{B})$ for some $\omega \in \mathcal{P}(\mathcal{B})$,³⁰² then $\omega(t)$ and ω lie in the same sector and hence in the same symplectic leaf of $\mathcal{P}(\mathcal{B})$.

To illustrate these concepts, let us return to our continuous field of C^* -algebras $\mathcal{A}^{(c)}$; cf. (2). It may not come as a great surprise that the canonical C^* -algebraic transition probabilities (27) on the pure state space of each fiber algebra $\mathcal{A}_{1/N}^{(c)}$ for $N < \infty$ converge to the classical transition probabilities (31) on the commutative limit algebra $\mathcal{A}_0^{(c)}$. Similarly, the C^* -algebraic Poisson structure (38) on each $\mathcal{P}(\mathcal{A}_{1/N}^{(c)})$ converges to zero. However, we know from the limit $\hbar \rightarrow 0$ of quantum mechanics that in generating classical behaviour on the limit algebra of a continuous field of C^* -algebras one should rescale the commutators; see Subsection 4.3

³⁰⁰All this can be boosted into an axiomatic structure into which both classical and quantum theory fit; see [Landsman, 1997; 1998].

³⁰¹In quantum theory the flow is defined for any t . In classical dynamics, (41) holds for all t for which $\rho(t)$ and $\sigma(t)$ are defined, cf. footnote 253.

³⁰²This assumption means that there exists a unitary representation $t \mapsto U_t$ of \mathbb{R} on \mathcal{H}_ω such that $\pi_\omega(\alpha_t(A)) = U_t \pi_\omega(A) U_t^*$ for all $A \in \mathcal{B}$ and all $t \in \mathbb{R}$.

and Section 5. Thus we replace the Poisson bracket (38) for $\mathcal{A}_{1/N}^{(c)}$ by

$$(42) \quad \{f, g\}(\omega) = iN\omega([df_\omega, dg_\omega]).$$

Thus rescaled, the Poisson brackets on the spaces $\mathcal{P}(\mathcal{A}_{1/N}^{(c)})$ turn out to converge to the canonical Poisson bracket (38) on $\mathcal{P}(\mathcal{A}_0^{(c)}) = \mathcal{S}(\mathcal{A}_1)$, instead of the zero bracket expected from the commutative nature of the limit algebra $\mathcal{A}_0^{(c)}$. Consequently, the symplectic leaves of the *full* state space $\mathcal{S}(\mathcal{A}_1)$ of the fiber algebra $\mathcal{A}_1^{(c)}$ become the symplectic leaves of the *pure* state space $\mathcal{S}(\mathcal{A}_1)$ of the fiber algebra $\mathcal{A}_0^{(c)}$. This is undoubtedly indicative of the origin of classical phase spaces and their Poisson structures in quantum theory.

More precisely, we have the following result [Duffield and Werner, 1992a]:

If $A = (A_0, A_1, A_2, \dots)$ and $A' = (A'_0, A'_1, A'_2, \dots)$ are continuous sections of $\mathcal{A}^{(c)}$ defined by symmetric sequences,³⁰³ then the sequence

$$(43) \quad (\{A_0, A'_0\}, i[A_1, A'_1], \dots, iN[A_N, A'_N], \dots)$$

defines a continuous section of $\mathcal{A}^{(c)}$.

This follows from an easy computation. In other words, although the sequence of commutators $[A_N, A'_N]$ converges to zero, the rescaled commutators $iN[A_N, A'_N] \in \mathcal{A}_N$ converge to the macroscopic observable $\{A_0, A'_0\} \in \mathcal{A}_0^{(c)} = C(\mathcal{S}(\mathcal{A}_1))$. Although it might seem perverse to reinterpret this result on the classical limit of a large quantum system in terms of quantization (which is the *opposite* of taking the classical limit), it is formally possible to do so (cf. Section 4.3) if we put

$$(44) \quad \hbar = \frac{1}{N}.$$

Using the axiom of choice if necessary, we devise a procedure that assigns a continuous section $A = (A_0, A_1, A_2, \dots)$ of our field to a given function $A_0 \in \mathcal{A}_0^{(c)}$. We write this as $A_N = \mathcal{Q}_{\frac{1}{N}}(A_0)$, and similarly $A'_N = \mathcal{Q}_{\frac{1}{N}}(A'_0)$. This choice need not be such that the sequence (43) is assigned to $\{A_0, A'_0\}$, but since the latter is the unique limit of (43), it must be that

$$(45) \quad \lim_{N \rightarrow \infty} \left\| iN \left[\mathcal{Q}_{\frac{1}{N}}(A_0), \mathcal{Q}_{\frac{1}{N}}(A'_0) \right] - \mathcal{Q}_{\frac{1}{N}}(\{A_0, A'_0\}) \right\| = 0.$$

Also note that (27) is just (12). Consequently (cf. (25) and surrounding text):

³⁰³The result does not hold for all continuous sections (i.e. for all approximately symmetric sequences), since, for example, the limiting functions A_0 and A'_0 may not be differentiable, so that their Poisson bracket does not exist. This problem occurs in all examples of deformation quantization. However, the class of sequences for which the claim is valid is larger than the symmetric ones alone. A sufficient condition on A and B for (43) to make sense is that $A_N = \sum_{M \leq N} j_{NM}(A_M^{(N)})$ (with $A_M^{(N)} \in \mathcal{A}_1^M$), such that $\lim_{N \rightarrow \infty} A_M^{(N)}$ exists (in norm) and $\sum_{M=1}^{\infty} M \sup_{N \geq M} \{\|A_M^{(N)}\|\} < \infty$. See [Duffield and Werner, 1992a].

The continuous field of C^* -algebras $\mathcal{A}^{(c)}$ defined by (2) and approximately symmetric sequences (and their limits) as continuous sections yields a deformation quantization of the phase space $\mathcal{S}(\mathcal{A}_1)$ (equipped with the Poisson bracket (38)) for any quantization map \mathcal{Q} .

For the dynamics this implies:

Let $H = (H_0, H_1, H_2, \dots)$ be a continuous section of $\mathcal{A}^{(c)}$ defined by a symmetric sequence,³⁰⁴ and let $A = (A_0, A_1, A_2, \dots)$ be an arbitrary continuous section of $\mathcal{A}^{(c)}$ (i.e. an approximately symmetric sequence). Then the sequence

$$(46) \quad (A_0(t), e^{iH_1 t} A_1 e^{-iH_1 t}, \dots, e^{iNH_N t} A_N e^{-iNH_N t}, \dots),$$

where $A_0(t)$ is the solution of the equations of motion with classical Hamiltonian H_0 ,³⁰⁵ defines a continuous section of $\mathcal{A}^{(c)}$.

In other words, for bounded symmetric sequences of Hamiltonians H_N the quantum dynamics restricted to macroscopic observables converges to the classical dynamics with Hamiltonian H_0 . Compare the positions of \hbar and N in (12) and (46), respectively, and rejoice in the reconfirmation of (44).

In contrast, the quasilocal observables are *not* well behaved as far as the $N \rightarrow \infty$ limit of the dynamics defined by such Hamiltonians is concerned. Namely, if (A_0, A_1, \dots) is a section of the continuous field $\mathcal{A}^{(q)}$, and (H_1, H_2, \dots) is any bounded symmetric sequence of Hamiltonians, then the sequence

$$(e^{iH_1 t} A_1 e^{-iH_1 t}, \dots, e^{iNH_N t} A_N e^{-iNH_N t}, \dots)$$

has no limit for $N \rightarrow \infty$, in that it cannot be extended by some $A_0(t)$ to a continuous section of $\mathcal{A}^{(q)}$. Indeed, this was the very reason why macroscopic observables were originally introduced in this context [Rieckers, 1984; Morchio and Strocchi, 1987; Bona, 1988; Unnerstall, 1990a; Raggio and Werner, 1989; Duffield and Werner, 1992a]. Instead, the natural finite- N Hamiltonians for which the limit $N \rightarrow \infty$ of the time-evolution on \mathcal{A}_1^N exists as a one-parameter automorphism group on $\mathcal{A}^{(q)}$ satisfy an appropriate locality condition, which excludes the global averages defining symmetric sequences.

6.6 Epilogue: Macroscopic observables and the measurement problem

In a renowned paper, Hepp [1972] suggested that macroscopic observables and superselection rules should play a role in the solution of the measurement problem of

³⁰⁴Once again, the result in fact holds for a larger class of Hamiltonians, namely the ones satisfying the conditions specified in footnote 303 [Duffield and Werner, 1992a]. The assumption that each Hamiltonian H_N lies in \mathcal{A}_1^N and hence is bounded is natural in lattice models, but is undesirable in general.

³⁰⁵See (18) and surrounding text.

quantum mechanics. He assumed that a macroscopic apparatus may be idealized as an infinite quantum system, whose algebra of observables \mathcal{A}_A has disjoint pure states. Referring to our discussion in Subsection 2.5 for context and notation, Hepp's basic idea (for which he claimed no originality) was that as a consequence of the measurement process the initial state vector $\Omega_I = \sum_n c_n \Psi_n \otimes I$ of system plus apparatus evolves into a final state vector $\Omega_F = \sum_n c_n \Psi_n \otimes \Phi_n$, in which each Φ_n lies in a different superselection sector of the Hilbert space of the apparatus (in other words, the corresponding states φ_n on \mathcal{A}_A are mutually disjoint). Consequently, although the initial state ω_I is pure, the final state ω_F is mixed. Moreover, because of the disjointness of the ω_n the final state ω_F has a unique decomposition $\omega_F = \sum_n |c_n|^2 \psi_n \otimes \varphi_n$ into pure states, and therefore admits a bona fide ignorance interpretation. Hepp therefore claimed with some justification that the measurement "reduces the wave packet", as desired in quantum measurement theory.

Even apart from the usual conceptual problem of passing from the collective of all terms in the final mixture to one actual measurement outcome, Hepp himself indicated a serious mathematical problem with this program. Namely, if the initial state is pure it must lie in a certain superselection sector (or equivalence class of states); but then the final state must lie in the very same sector if the time-evolution is Hamiltonian, or, more generally, automorphic (as we have seen in the preceding subsection). Alternatively, it follows from a more general lemma Hepp [1972] himself proved:

If two states ρ, σ on a C^ -algebra \mathcal{B} are disjoint and $\alpha : \mathcal{B} \rightarrow \mathcal{B}$ is an automorphism of \mathcal{B} , then $\rho \circ \alpha$ and $\sigma \circ \alpha$ are disjoint, too.*

To reach the negative conclusion above, one takes \mathcal{B} to be the algebra of observables of system and apparatus jointly, and computes back in time by choosing $\alpha = \alpha_{t_F - t_I}^{-1}$, where α_t is the one-parameter automorphism group on \mathcal{B} describing the joint time-evolution of system and apparatus (and t_I and t_F are the initial and final times of the measurement, respectively). However, Hepp pointed out that this conclusion may be circumvented if one admits the possibility that a measurement takes infinitely long to complete. For the limit $A \mapsto \lim_{t \rightarrow \infty} \alpha_t(A)$ (provided it exists in a suitable sense, e.g., weakly) does not necessarily yield an automorphism of \mathcal{B} . Hence a state — evolving in the Schrödinger picture by $\omega_t(A) \equiv \omega(\alpha_t(A))$ — may leave its sector in infinite time, a possibility Hepp actually demonstrated in a range of models; see also [Frigerio, 1974; Whitten-Wolfe and Emch, 1976; Araki, 1980; Bona, 1980; Hannabuss, 1984; Bub, 1988; Landsman, 1991; Frasca, 2003; 2004], and many other papers.

Despite the criticism that has been raised against the conclusion that a quantum-mechanical measurement requires an infinite apparatus and must take infinite time [Bell, 1975; Robinson, 1994; Landsman, 1995], and despite the fact that this procedure is quite against the spirit of von Neumann [1932], in whose widely accepted description measurements are practically instantaneous, this conclusion resonates well with the modern idea that quantum theory is universally valid and the classical

world has no absolute existence; cf. the Introduction. Furthermore, a quantum-mechanical measurement is nothing but a specific interaction, comparable with a scattering process; and it is quite uncontroversial that such a process takes infinite time to complete. Indeed, what would it mean for scattering to be over after some finite time? Which time? As we shall see in the next section, the theory of decoherence requires the limit $t \rightarrow \infty$ as well, and largely for the same mathematical reasons. There as well as in Hepp's approach, the limiting behaviour actually tends to be approached very quickly (on the pertinent time scale), and one needs to let $t \rightarrow \infty$ merely to make terms $\sim \exp -\gamma t$ (with $\gamma > 0$) zero rather than just very small. See also [Primas, 1997] for a less pragmatic point of view on the significance of this limit.

A more serious problem with Hepp's approach lies in his assumption that the time-evolution on the quasilocal algebra of observables of the infinite measurement apparatus (which in our class of examples would be $\mathcal{A}_0^{(q)}$) is automorphic. This, however, is by no means always the case; cf. the references listed near the end of Subsection 6.5. As we have seen, for certain natural Hamiltonian (and hence automorphic) time-evolutions at finite N the dynamics *has no limit* $N \rightarrow \infty$ on the algebra of quasilocal observables — let alone an automorphic one.

Nonetheless, Hepp's conclusion remains valid if we use the algebra $\mathcal{A}_0^{(c)}$ of macroscopic observables, on which (under suitable assumptions — see Subsection 6.5) Hamiltonian time-evolution on \mathcal{A}_1^N *does* have a limit as $N \rightarrow \infty$. For, as pointed out in Subsection 6.3, each superselection sector of $\mathcal{A}_0^{(q)}$ defines and is defined by a pure state on \mathcal{A}_1 , which in turn defines a sector of $\mathcal{A}_0^{(c)}$. Now the latter sector is simply a point in the pure state space $\mathcal{S}(\mathcal{A}_1)$ of the commutative C^* -algebra $\mathcal{A}_0^{(c)}$, so that Hepp's lemma quoted above boils down to the claim that if $\rho \neq \sigma$, then $\rho \circ \alpha \neq \sigma \circ \alpha$ for any automorphism α . This, of course, is a trivial property of any Hamiltonian time-evolution, and it follows once again that a transition from a pure pre-measurement state to a mixed post-measurement state on $\mathcal{A}_0^{(c)}$ is impossible in finite time. To avoid this conclusion, one should simply avoid the limit $N \rightarrow \infty$, which is the root of the $t \rightarrow \infty$ limit; see [Janssens, 2005].

What, then, does all this formalism mean for Schrödinger's cat? In our opinion, it confirms the impression that the appearance of a paradox rests upon an equivocation. Indeed, the problem arises because one oscillates between two mutually exclusive interpretations.³⁰⁶

Either one is a bohemian theorist who, in vacant or in pensive mood, puts off his or her glasses and merely contemplates whether the cat is dead or alive. Such a person studies the cat exclusively from the point of view of its macroscopic observables, so that he or she has to use a post-measurement state $\omega_F^{(c)}$ on the algebra $\mathcal{A}_0^{(c)}$. If $\omega_F^{(c)}$ is pure, it lies in $\mathcal{P}(\mathcal{A}_1)$ (unless the pre-measurement state was mixed). Such a state corresponds to a single superselection sector $[\omega_F^{(q)}]$ of $\mathcal{A}_0^{(q)}$, so that the cat is dead or alive. If, on the other hand, $\omega_F^{(c)}$ is mixed (which

³⁰⁶Does *complementarity* re-enter through the back door?

is what occurs if Schrödinger has his way), there is no problem in the first place: at the level of macroscopic observables one merely has a statistical description of the cat.

Or one is a hard-working experimental physicist of formidable power, who investigates the detailed microscopic constitution of the cat. For him or her the cat is always in a pure state on \mathcal{A}_1^N for some large N . This time the issue of life and death is not a matter of lazy observation and conclusion, but one of sheer endless experimentation and computation. From the point of view of such an observer, nothing is wrong with the cat being in a coherent superposition of two states that are actually quite close to each other microscopically — at least for the time being.

Either way, *the riddle does not exist* (Wittgenstein, TLP, §6.5).

7 WHY CLASSICAL STATES AND OBSERVABLES?

‘We have found a strange footprint on the shores of the unknown. We have devised profound theories, one after another, to account for its origins. At last, we have succeeded in reconstructing the creature that made the footprint. And lo! It is our own.’ [Eddington, 1920, 200–201]

The conclusion of Sections 5 and 6 is that quantum theory may give rise to classical behaviour *in certain states* and *with respect to certain observables*. For example, we have seen that in the limit $\hbar \rightarrow 0$ coherent states and operators of the form $\mathcal{Q}_\hbar(f)$, respectively, are appropriate, whereas in the limit $N \rightarrow \infty$ one should use classical states (*nomen est omen!*) as defined in Subsection 6.2 and macroscopic observables. If, instead, one uses superpositions of such states, or observables with the wrong limiting behaviour, no classical physics emerges. Thus the question remains why the world at large should happen to be in such states, and why we turn out to study this world with respect to the observables in question. This question found its original incarnation in the measurement problem (cf. Subsection 2.5), but this problem is really a figure-head for a much wider difficulty.

Over the last 25 years,³⁰⁷ two profound and original answers to this question have been proposed.

7.1 Decoherence

The first goes under the name of *decoherence*. Pioneering papers include [van Kampen, 1954; Zeh, 1970; Zurek, 1981; 1982],³⁰⁸ and [Joos and Zeh, 1985], and some recent reviews are [Bub, 1999; Auletta, 2001; Joos *et al.*, 2003; Zurek, 2003; Blanchard and Olkiewicz, 2003; Bacciagaluppi, 2004] and [Schlosshauer, 2004].³⁰⁹

³⁰⁷Though some say the basic idea of decoherence goes back to Heisenberg and Ludwig.

³⁰⁸See also [Zurek, 1991] and the subsequent debate in *Physics Today* [Zurek, 1993], which drew wide attention to decoherence.

³⁰⁹The website http://almaak.usc.edu/~tbrun/Data/decoherence_list.html contains an extensive list of references on decoherence.

More references will be given in due course. The existence (and excellence) of these reviews obviates the need for a detailed treatment of decoherence in this article, all the more so since at the time of writing this approach appears to be in a transitional stage, conceptually as well as mathematically (as will be evident from what follows). Thus we depart from the layout of our earlier chapters and restrict ourselves to a few personal comments.

1. Mathematically, decoherence boils down to the idea of adding one more link to the von Neumann chain (see Subsection 2.5) beyond $S + A$ (i.e. the system and the apparatus). Conceptually, however, there is a major difference between decoherence and older approaches that took such a step: whereas previously (e.g., in the hands of von Neumann, London and Bauer, Wigner, etc.)³¹⁰ the chain *converged towards the observer*, in decoherence it *diverges away from the observer*. Namely, the third and final link is now taken to be the *environment* (taken in a fairly literal sense in agreement with the intuitive meaning of the word). In particular, in realistic models the environment is treated as an infinite system (necessitating the limit $N \rightarrow \infty$), which has the consequence that (in simple models where the pointer has discrete spectrum) the post-measurement state $\sum_n c_n \Psi_n \otimes \Phi_n \otimes \chi_n$ (in which the χ_n are mutually orthogonal) is only reached in the limit $t \rightarrow \infty$. However, as already mentioned in Subsection 6.6, infinite time is only needed mathematically in order to make terms of the type $\sim \exp -\gamma t$ (with $\gamma > 0$) zero rather than just very small: in many models the inner products (χ_n, χ_m) are actually negligible for $n \neq m$ within surprisingly short time scales.³¹¹

If only in view of the need for limits of the type $N \rightarrow \infty$ (for the environment) and $t \rightarrow \infty$, in our opinion decoherence is best linked to stance 1 of the Introduction: its goal is to explain the approximate appearance of the classical world from quantum mechanics seen as a universally valid theory. However, decoherence has been claimed to support almost any opinion on the foundations of quantum mechanics; cf. [Bacciagaluppi, 2004] and [Schlosshauer, 2004] for a critical overview and also see Point 3 below.

2. Originally, decoherence entered the scene as a proposed solution to the measurement problem (in the precise form stated at the end of Subsection 2.5). For the restriction of the state $\sum_n c_n \Psi_n \otimes \Phi_n \otimes \chi_n$ to $S + A$ (i.e. its trace over the degrees of freedom of the environment) is mixed in the limit $t \rightarrow \infty$, which means that the quantum-mechanical interference between the states $\Psi_n \otimes \Phi_n$ for different values of n has become ‘delocalized’ to the environment, and accordingly is irrelevant if the latter is not observed (i.e. omitted from the description). Unfortunately, the application of the ignorance interpretation of the mixed post-measurement state of $S + A$ is illegal even from the point of view of stance 1 of the Introduction. The ignorance interpretation

³¹⁰See [Wheeler and Zurek, 1983].

³¹¹Cf. Tables 3.1 and 3.2 on pp. 66–67 of [Joos *et al.*, 2003].

is only valid if the environment is kept within the description *and* is classical (in having a commutative C^* -algebra of observables). The latter assumption [Primas, 1983], however, makes the decoherence solution to the measurement problem circular.³¹²

In fact, as quite rightly pointed out by Bacciagaluppi [2004], decoherence actually *aggravates* the measurement problem. Where previously this problem was believed to be man-made and relevant only to rather unusual laboratory situations (important as these might be for the foundations of physics), it has now become clear that “measurement” of a quantum system *by the environment* (instead of by an experimental physicist) happens everywhere and all the time: hence it remains even more miraculous than before that there is a single outcome after each such measurement. Thus decoherence as such does not provide a solution to the measurement problem [Leggett, 2002];³¹³ Adler, 2003; Joos and Zeh, 2003], but is in actual fact parasitic on such a solution.

3. There have been various responses to this insight. The dominant one has been to combine decoherence with some interpretation of quantum mechanics: decoherence then finds a home, while conversely the interpretation in question is usually enhanced by decoherence. In this context, the most popular of these has been the many-worlds interpretation, which, after decades of obscurity and derision, suddenly started to be greeted with a flourish of trumpets in the wake of the popularity of decoherence. See, for example, [Saunders, 1993; 1995; Joos *et al.*, 2003] and [Zurek, 2003]. In quantum cosmology circles, the consistent histories approach has been a popular partner to decoherence, often in combination with many worlds; see below. The importance of decoherence in the modal interpretation has been emphasized by Dieks [1989b] and Bene and Dieks [2002], and practically all authors on decoherence find the opportunity to pay some lip-service to Bohr in one way or another. See [Bacciagaluppi, 2004] and [Schlosshauer, 2004] for a critical assessment of all these combinations.

In our opinion, none of the established interpretations of quantum mechanics will do the job, leaving room for genuinely new ideas. One such idea is the *return of the environment*: instead of “tracing it out”, as in the original setting of decoherence theory, the environment should *not* be ignored! The

³¹²On the other hand, treating the environment *as if* it were classical might be an improvement on the Copenhagen ideology of treating the measurement apparatus *as if* it were classical (cf. Section 3).

³¹³In fact, Leggett’s argument only applies to strawman 3 of the Introduction and loses its force against stance 1. For his argument is that decoherence just removes the *evidence* for a given state (of Schrödinger’s cat type) to be a superposition, and accuses those claiming that this solves the measurement problem of committing the logical fallacy that removal of the evidence for a crime would undo the crime. But according to stance 1 *the crime is only defined relative to the evidence!* Leggett is quite right, however, in insisting on the ‘from “and” to “or” problem’ mentioned at the end of the Introduction.

essence of measurement has now been recognized to be the *redundancy* of the outcome (or “record”) of the measurement in the environment. It is this very redundancy of information about the underlying quantum object that “objectifies” it, in that the information becomes accessible to a large number of observers without necessarily disturbing the object³¹⁴ [Zurek, 2003; Olivier *et al.* 2004; Blume-Kohout and Zurek, 2004; 2005]. This insight (called “Quantum Darwinism”) has given rise to the “existential” interpretation of quantum mechanics due to Zurek [2003].

4. Another response to the failure of decoherence (and indeed all other approaches) to solve the measurement problem (in the sense of failing to win a general consensus) has been of a somewhat more pessimistic (or, some would say, pragmatic) kind: all attempts to explain the quantum world are given up, yielding to the point of view that ‘the appropriate aim of physics at the fundamental level then becomes the representation and manipulation of information’ [Bub, 2004]. Here ‘measuring instruments ultimately remain black boxes at some level’, and one concludes that all efforts to *understand* measurement (or, for that matter, EPR-correlations) are futile and pointless.³¹⁵
5. Night thoughts of a quantum physicist, then?³¹⁶ Not quite. Turning vice into virtue: rather than solving the measurement problem, the true significance of the decoherence program is that it gives conditions under which there is no measurement problem! Namely, foregoing an explanation of the transition from the state $\sum_n c_n \Psi_n \otimes \Phi_n \otimes \chi_n$ of $S + A + \mathcal{E}$ to a single one of the states $\Psi_n \otimes \Phi_n$ of $S + A$, at the heart of decoherence is the claim that each of the latter states is *robust* against coupling to the environment (provided the Hamiltonian is such that $\Psi_n \otimes \Phi_n$ tensored with some initial state $\mathcal{I}_{\mathcal{E}}$ of the environment indeed evolves into $\Psi_n \otimes \Phi_n \otimes \chi_n$, as assumed so far). This implies that each state $\Psi_n \otimes \Phi_n$ remains pure after coupling to the environment and subsequent restriction to the original system plus apparatus, so that at the end of the day the environment has had no influence on it. In other words, the real point of decoherence is the phenomenon of *einselection* (for *environment-induced superselection*), where a state is ‘einselected’ precisely when (given some interaction Hamiltonian) it possesses the stability property just mentioned. The claim, then, is that einselected states are often classical, or at least that classical states (in the sense mentioned at the beginning of this section) are classical precisely because they are robust against coupling to the environment. Provided this scenario indeed gives

³¹⁴Such objectification is claimed to yield an ‘operational definition of existence’ [Zurek, 2003, 749].

³¹⁵It is indeed in describing the transformation of quantum information (or entropy) to classical information during measurement that decoherence comes to its own and exhibits some of its greatest strength. Perhaps for this reason such thinking pervades also Zurek [2003].

³¹⁶[Kent, 2000]. Pun on the title of McCormmach [1982].

rise to the classical world (which remains to be shown in detail), it gives a dynamical explanation of it. But even short of having achieved this goal, the importance of the notion of einselection cannot be overstated; in our opinion, it is the most important and powerful idea in quantum theory since entanglement (which einselection, of course, attempts to undo!).

6. The measurement problem, and the associated distinction between system and apparatus on the one hand and environment on the other, can now be *omitted* from decoherence theory. Continuing the discussion in Subsection 3.4, the goal of decoherence should simply be to find the robust or einselected states of a object \mathcal{O} coupled to an environment \mathcal{E} , as well as the induced dynamics thereof (given the time-evolution of $\mathcal{O} + \mathcal{E}$). This search, however, must include the correct *identification* of the object \mathcal{O} within the total $\mathcal{S} + \mathcal{E}$, namely as a subsystem that actually *has* such robust states. Thus the Copenhagen idea that the Heisenberg cut between object and apparatus be movable (cf. Subsection 3.2) will not, in general, extend to the “Primas-Zurek” cut between object and environment. In traditional physics terminology, the problem is to find the right “dressing” of a quantum system so as to make at least some of its states robust against coupling to its environment [Amann and Primas, 1997; Brun and Hartle, 1999; Omnès, 2002]. In other words: *What is a system?* To mark this change in perspective, we now change notation from \mathcal{O} (for “object”) to \mathcal{S} (for “system”). Various tools for the solution of this problem within the decoherence program have now been developed — with increasing refinement and also increasing reliance on concepts from information theory [Zurek, 2003] — but the right setting for it seems the formalism of consistent histories, see below.
7. Various dynamical regimes haven been unearthed, each of which leads to a different class of robust states [Joos *et al.*, 2003; Zurek, 2003; Schlosshauer, 2004]. Here $H_{\mathcal{S}}$ is the system Hamiltonian, H_I is the interaction Hamiltonian between system and environment, and $H_{\mathcal{E}}$ is the environment Hamiltonian. As stated, no reference to measurement, object or apparatus need be made here.
 - In the regime $H_{\mathcal{S}} \ll H_I$, for suitable Hamiltonians the robust states are the traditional pointer states of quantum measurement theory. This regime conforms to von Neumann’s [1932] idea that quantum measurements be almost instantaneous. If, moreover, $H_{\mathcal{E}} \ll H_I$ as well — with or without a measurement context — then the decoherence mechanism turns out to be universal in being independent of the details of \mathcal{E} and $H_{\mathcal{E}}$ [Strunz *et al.*, 2003].
 - If $H_{\mathcal{S}} \approx H_I$, then (at least in models of quantum Brownian motion) the robust states are coherent states (either of the traditional Schrödinger type, or of a more general nature as defined in Subsection 5.1); see [Zurek *et al.*, 1993] and [Zurek, 2003]. This case is, of course, of supreme

importance for the physical relevance of the results quoted in our Section 5 above, and — if only for this reason — decoherence theory would benefit from more interaction with mathematically rigorous results on quantum stochastic analysis.³¹⁷

- Finally, if $H_S \gg H_I$, then the robust states turn out to be eigenstates of the system Hamiltonian H_S [Paz and Zurek, 1999; Ollivier *et al.*, 2004]. In view of our discussion of such states in Subsections 5.5 and 5.6, this shows that robust states are not necessarily classical. It should be mentioned that in this context decoherence theory largely coincides with standard atomic physics, in which the atom is taken to be the system \mathcal{S} and the radiation field plays the role of the environment \mathcal{E} ; see [Gustafson and Sigal, 2003] for a mathematically minded introductory treatment and [Bach *et al.*, 1998; 1999] for a full (mathematical) meal.
8. Further to the above clarification of the role of energy eigenstates, decoherence also has had important things to say about quantum chaos [Zurek, 2003; Joos *et al.*, 2003]. Referring to our discussion of wave packet revival in Subsection 2.4, we have seen that in atomic physics wave packets do not behave classically on long time scales. Perhaps surprisingly, this is even true for certain chaotic macroscopic systems: cf. the case of Hyperion mentioned in the Introduction and at the end of Subsection 5.2. Decoherence now replaces the underlying superposition by a classical probability distribution, which reflects the chaotic nature of the limiting classical dynamics. Once again, the transition from the pertinent pure state of system plus environment to a *single* observed system state remains clouded in mystery. But granted this transition, decoherence sheds new light on classical chaos and circumvents at least the most flagrant clashes with observation.³¹⁸
 9. Robustness and einselection form the state side or Schrödinger picture of decoherence. Of course, there should also be a corresponding observable side or Heisenberg picture of decoherence. But the transition between the two pictures is more subtle than in the quantum mechanics of closed systems. In the Schrödinger picture, the whole point of einselection is that most pure states simply disappear from the scene. This may be beautifully visualized on the example of a two-level system with Hilbert space $\mathcal{H}_S = \mathbb{C}^2$ [Zurek, 2003]. If \uparrow and \downarrow (cf. (33)) happen to be the robust vector states of the system after coupling to an appropriate environment, and if we identify the corresponding density matrices with the north-pole $(0, 0, 1) \in B^3$ and the south-pole $(0, 0, -1) \in B^3$, respectively (cf. (3)), then following decoherence all other states move towards the axis connecting the north- and south poles

³¹⁷Cf. [Davies, 1976; Accardi *et al.*, 1990; Parthasarathy, 1992; Streater, 2000; Kümmerer, 2002; Maassen, 2003], etc.

³¹⁸It should be mentioned, though, that any successful mechanism explaining the transition from quantum to classical should have this feature, so that at the end of the day decoherence might turn out to be a red herring here.

(i.e. the intersection of the z -axis with B^3) as $t \rightarrow \infty$. In the Heisenberg picture, this disappearance of all pure states except two corresponds to the reduction of the full algebra of observables $M_2(\mathbb{C})$ of the system to its diagonal (and hence commutative) subalgebra $\mathbb{C} \oplus \mathbb{C}$ in the same limit. For it is only the latter algebra that contains enough elements to distinguish \uparrow and \downarrow without containing observables detecting interference terms between these pure states.

10. To understand this in a more abstract and general way, we recall the mathematical relationship between pure states and observables [Landsman, 1998]. The passage from a C^* -algebra \mathcal{A} of observables of a given system to its pure states is well known: as a set, the pure state space $\mathcal{P}(\mathcal{A})$ is the extreme boundary of the total state space $\mathcal{S}(\mathcal{A})$ (cf. footnote 259). In order to reconstruct \mathcal{A} from $\mathcal{P}(\mathcal{A})$, the latter needs to be equipped with the structure of a transition probability space (see Subsection 6.3) through (27). Each element $A \in \mathcal{A}$ defines a function \hat{A} on $\mathcal{P}(\mathcal{A})$ by $\hat{A}(\omega) = \omega(A)$. Now, in the simple case that \mathcal{A} is finite-dimensional (and hence a direct sum of matrix algebras), one can show that each function \hat{A} is a finite linear combination of the form $\hat{A} = \sum_i p_{\omega_i}$, where $\omega_i \in \mathcal{P}(\mathcal{A})$ and the elementary functions p_ρ on $\mathcal{P}(\mathcal{A})$ are defined by $p_\rho(\sigma) = p(\rho, \sigma)$. Conversely, each such linear combination defines a function \hat{A} for some $A \in \mathcal{A}$. Thus the elements of \mathcal{A} (seen as functions on the pure state space $\mathcal{P}(\mathcal{A})$) are just the transition probabilities and linear combinations thereof. The algebraic structure of \mathcal{A} may then be reconstructed from the structure of $\mathcal{P}(\mathcal{A})$ as a Poisson space with a transition probability (cf. Subsection 6.5). In this sense $\mathcal{P}(\mathcal{A})$ uniquely determines the algebra of observables of which it is the pure state space. For example, the space consisting of two points with classical transition probabilities (31) leads to the commutative algebra $\mathcal{A} = \mathbb{C} \oplus \mathbb{C}$, whereas the unit two-sphere in \mathbb{R}^3 with transition probabilities (32) yields $\mathcal{A} = M_2(\mathbb{C})$.

This reconstruction procedure may be generalized to arbitrary C^* -algebras [Landsman, 1998], and defines the precise connection between the Schrödinger picture and the Heisenberg picture that is relevant to decoherence. These pictures are equivalent, but in practice the reconstruction procedure may be difficult to carry through.

11. For this reason it is of interest to have a direct description of decoherence in the Heisenberg picture. Such a description has been developed by Blanchard and Olkiewicz [2003], partly on the basis of earlier results by Olkiewicz [1999a,b; 2000]. Mathematically, their approach is more powerful than the Schrödinger picture on which most of the literature on decoherence is based. Let $\mathcal{A}_S = \mathcal{B}(\mathcal{H}_S)$ and $\mathcal{A}_E = \mathcal{B}(\mathcal{H}_E)$, and assume one has a total Hamiltonian H acting on $\mathcal{H}_S \otimes \mathcal{H}_E$ as well as a fixed state of the environment, represented by a density matrix ρ_E (often taken to be a thermal equilibrium state). If ρ_S is a density matrix on \mathcal{H}_S (so that the total state is $\rho_S \otimes \rho_E$), the Schrödinger

picture approach to decoherence (and more generally to the quantum theory of open systems) is based on the time-evolution

$$(1) \quad \rho_S(t) = \text{Tr}_{\mathcal{H}_E} \left(e^{-\frac{it}{\hbar}H} \rho_S \otimes \rho_E e^{\frac{it}{\hbar}H} \right).$$

The Heisenberg picture, on the other hand, is based on the associated operator time-evolution for $A \in \mathcal{B}(\mathcal{H}_S)$ given by

$$(2) \quad A(t) = \text{Tr}_{\mathcal{H}_E} \left(\rho_E e^{\frac{it}{\hbar}H} A \otimes 1 e^{-\frac{it}{\hbar}H} \right),$$

since this yields the equivalence of the Schrödinger and Heisenberg pictures expressed by

$$(3) \quad \text{Tr}_{\mathcal{H}_S} (\rho_S(t)A) = \text{Tr}_{\mathcal{H}_S} (\rho_S A(t)).$$

More generally, let \mathcal{A}_S and \mathcal{A}_E be unital C^* -algebras with spatial tensor product $\mathcal{A}_S \otimes \mathcal{A}_E$, equipped with a time-evolution α_t and a fixed state ω_E on \mathcal{A}_E . This defines a conditional expectation $P_E : \mathcal{A}_S \otimes \mathcal{A}_E \rightarrow \mathcal{A}_S$ by linear and continuous extension of $P_E(A \otimes B) = A\omega_E(B)$, and consequently a reduced time-evolution $A \mapsto A(t)$ on \mathcal{A}_S via

$$(4) \quad A(t) = P_E(\alpha_t(A \otimes 1)).$$

See, for example, Alicki and Lendi [1987]; in our context, this generality is crucial for the potential emergence of continuous classical phase spaces; see below.³¹⁹ Now the key point is that decoherence is described by a decomposition $\mathcal{A}_S = \mathcal{A}_S^{(1)} \oplus \mathcal{A}_S^{(2)}$ as a vector space (not as a C^* -algebra), where $\mathcal{A}_S^{(1)}$ is a C^* -algebra, with the property that $\lim_{t \rightarrow \infty} A(t) = 0$ (weakly) for all $A \in \mathcal{A}_S^{(2)}$, whereas $A \mapsto A(t)$ is an automorphism on $\mathcal{A}_S^{(1)}$ for each finite t . Consequently, $\mathcal{A}_S^{(1)}$ is the effective algebra of observables after decoherence, and it is precisely the pure states on $\mathcal{A}_S^{(1)}$ that are robust or einselected in the sense discussed before.

- 12. For example, if $\mathcal{A}_S = M_2(\mathbb{C})$ and the states \uparrow and \downarrow are robust under decoherence, then $\mathcal{A}_S^{(1)} = \mathbb{C} \oplus \mathbb{C}$ and $\mathcal{A}_S^{(2)}$ consists of all 2×2 matrices with zeros on the diagonal. In this example $\mathcal{A}_S^{(1)}$ is commutative hence classical, but this may not be the case in general. But if it is, the automorphic time-evolution on $\mathcal{A}_S^{(1)}$ induces a classical flow on its structure space, which should be shown to be Hamiltonian using the techniques of Section 6.³²⁰

³¹⁹For technical reasons Blanchard and Olkiewicz [2003] assume \mathcal{A}_S to be a von Neumann algebra with trivial center.

³²⁰Since on the assumption in the preceding footnote $\mathcal{A}_S^{(1)}$ is a commutative von Neumann algebra one should define the structure space in an indirect way; see [Blanchard and Olkiewicz, 2003].

In any case, there will be some sort of classical behaviour of the decohered system whenever $\mathcal{A}_S^{(1)}$ has a nontrivial center.³²¹ If this center is discrete, then the induced time-evolution on it is necessarily trivial, and one has the typical measurement situation where the center in question is generated by the projections on the eigenstates of a pointer observable with discrete spectrum. This is generic for the case where \mathcal{A}_S is a type I factor. However, type II and III factors may give rise to continuous classical systems with nontrivial time-evolution; see [Lugiewicz and Olkiewicz, 2002; 2003]. We cannot do justice here to the full technical details and complications involved here. But we would like to emphasize that further to quantum field theory and the theory of the thermodynamic limit, the present context of decoherence should provide important motivation for specialists in the foundations of quantum theory to learn the theory of operator algebras.³²²

7.2 Consistent histories

Whilst doing so, one is well advised to work even harder and simultaneously familiarize oneself with *consistent histories*. This approach to quantum theory was pioneered by Griffiths [1984] and was subsequently taken up by Omnès [1992] and others. Independently, Gell-Mann and Hartle [1990; 1993] proposed analogous ideas. Like decoherence, the consistent histories method has been the subject of lengthy reviews [Hartle, 1995] and even books [Omnès, 1994; 1999; Griffiths, 2002] by the founders. See also the reviews by Kiefer [2003] and Halliwell [2004], the critiques by Dowker and Kent [1996], Kent [1998], Bub [1999], and Bassi and Ghirardi [2000], as well as the various mathematical reformulations and reinterpretations of the consistent histories program [Isham, 1994; 1997; Isham and Linden, 1994; 1995; Isham *et al.*, 1994; Isham and Butterfield, 2000; Rudolph, 1996a; 1996b; 2000; Rudolph and Wright, 1999].

The relationship between consistent histories and decoherence is somewhat peculiar: on the one hand, decoherence is a natural mechanism through which appropriate sets of histories become (approximately) consistent, but on the other hand these approaches appear to have quite different points of departure. Namely, where decoherence starts from the idea that (quantum) systems are naturally coupled to their environments and therefore have to be treated as *open* systems, the aim of consistent histories is to deal with *closed* quantum systems such as the Universe, without a priori talking about measurements or observers. However, this distinction is merely historical: as we have seen in item 6 in the previous subsection, the dividing line between a system and its environment should be seen as a dynamical entity to be drawn according to certain stability criteria, so that even in decoherence theory one should really study the system plus its environment as a whole from the outset.³²³ And this is precisely what consistent historians do.

³²¹This is possible even when \mathcal{A}_S is a factor!

³²²See the references in footnote 7.

³²³This renders the distinction between “open” and “closed” systems a bit of a red herring, as

As in the preceding subsection, and for exactly the same reasons, we format our treatment of consistent histories as a list of items open to discussion.

1. The starting point of the consistent histories formulation of quantum theory is conventional: one has a Hilbert space \mathcal{H} , a state ρ , taken to be the initial state of the total system under consideration (realized as a density matrix on \mathcal{H}) and a Hamiltonian H (defined as a self-adjoint operator on \mathcal{H}). What is unconventional is that this total system may well be the entire Universe. Each property α of the total system is mathematically represented by a projection P_α on \mathcal{H} ; for example, if α is the property that the energy takes some value ϵ , then the operator P_α is the projection onto the associated eigenspace (assuming ϵ belongs to the discrete spectrum of H). In the Heisenberg picture, P_α evolves in time as $P_\alpha(t)$ according to (12); note that $P_\alpha(t)$ is once again a projection.

A *history* \mathbb{H}_A is a chain of properties (or propositions) $(\alpha_1(t_1), \dots, \alpha_n(t_n))$ indexed by n different times $t_1 < \dots < t_n$; here A is a multi-label incorporating both the properties $(\alpha_1, \dots, \alpha_n)$ and the times (t_1, \dots, t_n) . Such a history indicates that each property α_i holds at time t_i , $i = 1, \dots, n$. Such a history may be taken to be a collection $\{\alpha(t)\}_{t \in \mathbb{R}}$ defined for all times, but for simplicity one usually assumes that $\alpha(t) \neq 1$ (where 1 is the trivial property that always holds) only for a finite set of times t ; this set is precisely $\{t_1, \dots, t_n\}$. An example suggested by Heisenberg (1927) is to take α_i to be the property that a particle moving through a Wilson cloud chamber may be found in a cell $\Delta_i \subset \mathbb{R}^6$ of its phase space; the history $(\alpha_1(t_1), \dots, \alpha_n(t_n))$ then denotes the state of affairs in which the particle is in cell Δ_1 at time t_1 , subsequently is in cell Δ_2 at time t_2 , etcetera. Nothing is stated about the particle's behaviour at intermediate times. Another example of a history is provided by the double slit experiment, where α_1 is the particle's launch at the source at t_1 (which is usually omitted from the description), α_2 is the particle passing through (e.g.) the upper slit at t_2 , and α_3 is the detection of the particle at some location L at the screen at t_3 . As we all know, there is a potential problem with this history, which will be clarified below in the present framework.

The fundamental claim of the consistent historians seems to be that quantum theory should do no more (or less) than making predictions about the probabilities that histories occur. What these probabilities actually mean remains obscure (except perhaps when they are close to zero or one, or when reference is made to some measurement context; see [Hartle, 2005]), but let us first see when and how one can define them. The only potentially meaningful mathematical expression (within quantum mechanics) for the probability of

even in decoherence theory the totality of the system plus its environment is treated as a closed system.

a history \mathbb{H}_A with respect to a state ρ is [Groenewold, 1952; Wigner, 1963]

$$(5) \quad p(\mathbb{H}_A) = \text{Tr}(C_A \rho C_A^*),$$

where

$$(6) \quad C_A = P_{\alpha_n}(t_n) \cdots P_{\alpha_1}(t_1).$$

Note that C_A is generally not a projection (and hence a property) itself (unless all P_{α_i} mutually commute). In particular, when $\rho = [\Psi]$ is a pure state (defined by some unit vector $\Psi \in \mathcal{H}$), one simply has

$$(7) \quad p(\mathbb{H}_A) = \|C_A \Psi\|^2 = \|P_{\alpha_n}(t_n) \cdots P_{\alpha_1}(t_1) \Psi\|^2.$$

When $n = 1$ this just yields the Born rule. Conversely, see Isham (1994) for a derivation of (5) from the Born rule.³²⁴

2. Whatever one might think about the metaphysics of quantum mechanics, a probability makes no sense whatsoever when it is only attributed to a single history (except when it is exactly zero or one). The least one should have is something like a sample space (or event space) of histories, each (measurable) subset of which is assigned some probability such that the usual (Kolmogorov) rules are satisfied. This is a (well-known) problem even for a single time t and a single projection P_α (i.e. $n = 1$). In that case, the problem is solved by finding a self-adjoint operator A of which P_α is a spectral projection, so that the sample space is taken to be the spectrum $\sigma(A)$ of A , with $\alpha \subset \sigma(A)$. Given P_α , the choice of A is by no means unique, of course; different choices may lead to different and incompatible sample spaces. In practice, one usually starts from A and derives the P_α as its spectral projections $P_\alpha = \int_\alpha dP(\lambda)$, given that the spectral resolution of A is $A = \int_{\mathbb{R}} dP(\lambda) \lambda$. Subsequently, one may then either *coarse-grain* or *fine-grain* this sample space. The former is done by finding a partition $\sigma(A) = \coprod_i \alpha_i$ (disjoint union), and only admitting elements of the σ -algebra generated by the α_i as events (along with the associated spectral projection P_{α_i}), instead of all (measurable) subsets of $\sigma(A)$. To perform fine-graining, one supplements A by operators that commute with A as well as with each other, so that the new sample space is the joint spectrum of the ensuing family of mutually commuting operators.

In any case, in what follows it turns out to be convenient to work with the projections P_α instead of the subsets α of the sample space; the above discussion then amounts to extending the given projection on \mathcal{H} to some

³²⁴See also [Zurek, 2005] for a novel derivation of the Born rule, as well as the ensuing discussion in [Schlosshauer, 2004].

Boolean sublattice of the lattice $\mathcal{P}(\mathcal{H})$ of all projections on \mathcal{H} .³²⁵ Any state ρ then defines a probability measure on this sublattice in the usual way [Beltrametti and Cassinelli, 1984].

3. Generalizing this to the multi-time case is not a trivial task, somewhat facilitated by the following device (Isham, 1994). Put $\mathcal{H}^N = \otimes^N \mathcal{H}$, where N is the cardinality of the set of all times t_i relevant to the histories in the given collection,³²⁶ and, for a given history \mathbb{H}_A , define

$$(8) \quad \mathbb{C}_A = P_{\alpha_n}(t_n) \otimes \cdots \otimes P_{\alpha_1}(t_1).$$

Here $P_{\alpha_i}(t_i)$ acts on the copy of \mathcal{H} in the tensor product \mathcal{H}^N labeled by t_i , so to speak. Note that \mathbb{C}_A is a projection on \mathcal{H}^N (whereas C_A in (6) is generally *not* a projection on \mathcal{H}). Furthermore, given a density matrix ρ on \mathcal{H} as above, define the *decoherence functional* d as a map from pairs of histories into \mathbb{C} by

$$(9) \quad d(\mathbb{H}_A, \mathbb{H}_B) = \text{Tr}(C_A \rho C_B^*).$$

The main point of the consistent histories approach may now be summarized as follows: a collection $\{\mathbb{H}_A\}_{A \in \mathbb{A}}$ of histories can be regarded as a sample space on which a state ρ defines a probability measure via (5), which of course amounts to

$$(10) \quad p(\mathbb{H}_A) = d(\mathbb{H}_A, \mathbb{H}_A),$$

provided that:

- (a) The operators $\{\mathbb{C}_A\}_{A \in \mathbb{A}}$ form a Boolean sublattice of the lattice $\mathcal{P}(\mathcal{H}^N)$ of all projections on \mathcal{H}^N ;
- (b) The real part of $d(\mathbb{H}_A, \mathbb{H}_B)$ vanishes whenever \mathbb{H}_A is disjoint from \mathbb{H}_B .³²⁷

In that case, the set $\{\mathbb{H}_A\}_{A \in \mathbb{A}}$ is called *consistent*. It is important to realize that the possible consistency of a given set of histories depends (trivially) not only on this set, but in addition on the dynamics and on the initial state.

Consistent sets of histories generalize families of commuting projections at a single time. There is no great loss in replacing the second condition by the vanishing of $d(\mathbb{H}_A, \mathbb{H}_B)$ itself, in which case the histories \mathbb{H}_A and \mathbb{H}_B are

³²⁵This sublattice is supposed to the unit of $\mathcal{P}(\mathcal{H})$, i.e. the unit operator on \mathcal{H} , as well as the zero projection. This comment also applies to the Boolean sublattice of $\mathcal{P}(\mathcal{H}^N)$ discussed below.

³²⁶See the mathematical references above for the case $N = \infty$.

³²⁷This means that $\mathbb{C}_A \mathbb{C}_B = 0$; equivalently, $P_{\alpha_i}(t_i) P_{\beta_i}(t_i) = 0$ for at least one time t_i . This condition guarantees that the probability (10) is additive on disjoint histories.

said to *decohere*.³²⁸ For example, in the double slit experiment the pair of histories $\{\mathbb{H}_A, \mathbb{H}_B\}$ where $\alpha_1 = \beta_1$ is the particle's launch at the source at t_1 , α_2 (β_2) is the particle passing through the upper (lower) slit at t_2 , and $\alpha_3 = \beta_3$ is the detection of the particle at some location L at the screen, is *not consistent*. It becomes consistent, however, when the particle's passage through either one of the slits is recorded (or measured) *without the recording device being included in the histories* (if it is, nothing would be gained). This is reminiscent of the von Neumann chain in quantum measurement theory, which indeed provides an abstract setting for decoherence (cf. item 1 in the preceding subsection). Alternatively, the set can be made consistent by omitting α_2 and β_2 . See [Griffiths, 2002] for a more extensive discussion of the double slit experiment in the language of consistent histories.

More generally, coarse-graining by simply leaving out certain properties is often a promising attempt to make a given inconsistent set consistent; if the original history was already consistent, it can never become inconsistent by doing so. Fine-graining (by embedding into a larger set), on the other hand, is a dangerous act in that it may render a consistent set inconsistent.

4. What does it all mean? Each choice of a consistent set defines a “universe of discourse” within which one can apply classical probability theory and classical logic [Omnès, 1992]. In this sense the consistent historians are quite faithful to the Copenhagen spirit (as most of them acknowledge): *in order to understand it, the quantum world has to be looked at through classical glasses*. In our opinion, no convincing case has ever been made for the absolute necessity of this Bohrian stance (cf. Subsection 3.1), but accepting it, the consistent histories approach is superior to Copenhagen in not relying on measurement as an a priori ingredient in the interpretation of quantum mechanics.³²⁹ It is also more powerful than the decoherence approach in turning the notion of a system into a dynamical variable: different consistent sets describe different systems (and hence different environments, defined as the rest of the Universe); cf. item 6 in the previous subsection.³³⁰ In other words, the choice of a consistent set boils down to a choice of “relevant variables” against “irrelevant” ones omitted from the description. As indeed

³²⁸Consistent historians use this terminology in a different way from decoherence theorists. By definition, any two histories involving only a single time are consistent (or, indeed, “decohere”) iff condition (a) above holds; condition (b) is trivially satisfied in that case, and becomes relevant only when more than one time is considered. However, in decoherence theory the reduced density matrix at some given time does not trivially “decohere” at all; the whole point of the (original) decoherence program was to provide models in which this happens (if only approximately) because of the coupling of the system with its environment. Having said this, within the context of models there are close links between consistency (or decoherence) of multi-time histories and decoherence of reduced density matrices, as the former is often (approximately) achieved by the same kind of dynamical mechanisms that lead to the latter.

³²⁹See [Hartle, 2005] for an analysis of the connection between consistent histories and the Copenhagen interpretation and others.

³³⁰Technically, as the commutant of the projections occurring in a given history.

stressed in the literature, the act of identification of a certain consistent set as a universe of discourse is itself nothing but a coarse-graining of the Universe as a whole.

5. But these conceptual successes come with a price tag. Firstly, *consistent sets turn out not to exist in realistic models* (at least if the histories in the set carry more than one time variable). This has been recognized from the beginning of the program, the response being that one has to deal with approximately consistent sets for which (the real part of) $d(\mathbb{H}_A, \mathbb{H}_B)$ is merely very small. Furthermore, even the definition of a history often cannot be given in terms of projections. For example, in Heisenberg's cloud chamber example (see item 1 above), because of his very own uncertainty principle it is impossible to write down the corresponding projections P_{α_i} . A natural candidate would be $P_\alpha = \mathcal{Q}_\hbar^B(\chi_\Delta)$, cf. (19) and (28), but in view of (21) this operator fails to satisfy $P_\alpha^2 = P_\alpha$, so that it is not a projection (although it does satisfy the second defining property of a projection $P_\alpha^* = P_\alpha$). This merely reflects the usual property $\mathcal{Q}(f)^2 \neq \mathcal{Q}(f^2)$ of any quantization method, and necessitates the use of approximate projections [Omnès, 1997]. Indeed, this point calls for a reformulation of the entire consistent histories approach in terms of positive operators instead of projections [Rudolph, 1996a,b].

These are probably not serious problems; indeed, the recognition that classicality emerges from quantum theory only in an approximate sense (conceptually as well as mathematically) is a profound one (see the Introduction), and it rather should be counted among its blessings that the consistent histories program has so far confirmed it.

6. What is potentially more troubling is that consistency by no means implies classicality *beyond the ability (within a given consistent set) to assign classical probabilities and to use classical logic*. Quite to the contrary, neither Schrödinger cat states nor histories that look classical at each time but follow utterly unclassical trajectories in time are forbidden by the consistency conditions alone [Dowker and Kent, 1996]. But is this a genuine problem, except to those who still believe that the earth is at the centre of the Universe and/or that humans are privileged observers? It just seems to be the case that — at least according to the consistent historians — the ontological landscape laid out by quantum theory is far more “inhuman” (or some would say “obscure”) than the one we inherited from Bohr, in the sense that most consistent sets bear no obvious relationship to the world that *we* observe. In attempting to make sense of these, no appeal to “complementarity” will do now: for one, the complementary pictures of the quantum world called for by Bohr were classical in a much stronger sense than generic consistent sets are, and on top of that Bohr asked us to only think about two such pictures, as opposed to the innumerable consistent sets offered to us. Our conclusion is that, much as decoherence does not solve the measurement problem but rather aggravates it (see item 2 in the preceding subsection), also *consistent*

histories actually make the problem of interpreting quantum mechanics more difficult than it was thought to be before. In any case, it is beyond doubt that the consistent historians have significantly deepened our understanding of quantum theory — at the very least by providing a good bookkeeping device!

7. Considerable progress has been made in the task of identifying at least *some* (approximately) consistent sets that display (approximate) classical behaviour in the full sense of the word [Gell-Mann and Hartle, 1993; Omnès, 1992; 1997; Halliwell, 1998; 2000; 2004; Brun and Hartle, 1999; Bosse and Hartle, 2005]. Indeed, in our opinion studies of this type form the main concrete outcome of the consistent histories program. The idea is to find a consistent set $\{\mathbb{H}_A\}_{A \in \mathbb{A}}$ with three decisive properties:
 - (a) Its elements (i.e. histories) are strings of propositions with a classical interpretation;
 - (b) Any history in the set that delineates a classical trajectory (i.e. a solution of appropriate classical equations of motion) has probability (10) close to unity, and any history following a classically impossible trajectory has probability close to zero;
 - (c) The description is sufficiently coarse-grained to achieve consistency, but is sufficiently fine-grained to turn the deterministic equations of motion following from (b) into a closed system.

When these goals are met, it is in this sense (no more, no less) that the consistent histories program can claim with some justification that it has indicated (or even explained) ‘How the quantum Universe becomes classical’ [Halliwell, 2005].

Examples of propositions with a classical interpretation are quantized classical observables with a recognizable interpretation (such as the operators $Q_h^B(\chi_\Delta)$ mentioned in item 5), macroscopic observables of the kind studied in Subsection 6.1, and hydrodynamic variables (i.e. spatial integrals over conserved currents). These represent three different levels of classicality, which in principle are connected through mutual fine- or coarse-grainings.³³¹ The first are sufficiently coarse-grained to achieve consistency only in the limit $\hbar \rightarrow 0$ (cf. Section 5), whereas the latter two are already coarse-grained by their very nature. Even so, also the initial state will have to be “classical” in some sense in order to achieve the three targets (a) - (c).

All this is quite impressive, but we would like to state our opinion that neither decoherence nor consistent histories can stand on their own in explaining the appearance of the classical world. Promising as these approaches are, they have to

³³¹The study of these connections is relevant to the program laid out in this paper, but really belongs to classical physics *per se*; think of the derivation of the Navier–Stokes equations from Newton’s equations.

be combined at least with limiting techniques of the type described in Sections 5 and 6 — not to speak of the need for a new metaphysics! For even if it is granted that decoherence yields the disappearance of superpositions of Schrödinger cat type, or that consistent historians give us consistent sets none of whose elements contain such superpositions among their properties, this by no means suffices to explain the emergence of classical phase spaces and flows thereon determined by classical equations of motion. Since so far the approaches cited in Sections 5 and 6 have hardly been combined with the decoherence and/or the consistent histories program, a full explanation of the classical world from quantum theory is still in its infancy. This is not merely true at the technical level, but also conceptually; what has been done so far only represents a modest beginning. On the positive side, here lies an attractive challenge for mathematically minded researchers in the foundations of physics!

8 EPILOGUE

As a sobering closing note, one should not forget that whatever one's achievements in identifying a “classical realm” in quantum mechanics, the theory continues to incorporate another realm, the pure quantum world, that the young Heisenberg first gained access to, if not through his mathematics, then perhaps through the music of his favourite composer, Beethoven. This world beyond ken has never been better described than by Hoffmann [1810] in his essay on Beethoven's instrumental music, and we find it appropriate to end this paper by quoting at some length from it.³³²

Should one, whenever music is discussed as an independent art, not always be referred to instrumental music which, refusing the help of any other art (of poetry), expresses the unique essence of art that can only be recognized in it? It is the most romantic of all arts, one would almost want to say, the only truly romantic one, for only the infinite is its source. Orpheus' lyre opened the gates of the underworld. Music opens to man an unknown realm, a world that has nothing in common with the outer sensual world that surrounds him, a realm in which he leaves behind all of his feelings of certainty, in order to abandon himself to an unspeakable longing. (...)

Beethoven's instrumental music opens to us the realm of the gigantic and unfathomable. Glowing rays of light shoot through the dark night of this realm, and we see gigantic shadows swaying back and forth, encircling us closer and closer, destroying us (...) Beethoven's music moves the levers of fear, of shudder, of horror, of pain and thus awakens that infinite longing that is the essence of romanticism. Therefore, he is a purely romantic composer, and may it not be because of it, that

³³²Translation copyright: Ingrid Schwaegermann [2001].

to him, vocal music that does not allow for the character of infinite longing — but, through words, achieves certain effects, as they are not present in the realm of the infinite — is harder?(...)

What instrumental work of Beethoven confirms this to a higher degree than his magnificent and profound Symphony in c-Minor. Irresistibly, this wonderful composition leads its listeners in an increasing climax towards the realm of the spirits and the infinite.(...) Only that composer truly penetrates into the secrets of harmony who is able to have an effect on human emotions through them; to him, relationships of numbers, which, to the Grammarian, must remain dead and stiff mathematical examples without genius, are magic potions from which he lets a miraculous world emerge. (...)

Instrumental music, wherever it wants to only work through itself and not perhaps for a certain dramatic purpose, has to avoid all unimportant punning, all dallying. It seeks out the deep mind for premonitions of joy that, more beautiful and wonderful than those of this limited world, have come to us from an unknown country, and spark an inner, wonderful flame in our chests, a higher expression than mere words — that are only of this earth — can spark.

ACKNOWLEDGEMENTS

The author is indebted to Stephan de Bièvre, Jeremy Butterfield, Dennis Dieks, Jim Hartle, Gijs Tuynman, Steven Zelditch, and Wojciech Zurek for detailed comments on various drafts of this paper. The final version has greatly benefited from the 7 Pines Meeting on ‘The Classical-Quantum Borderland’ (May, 2005); the author wishes to express his gratitude to Lee Gohlike and the Board of the 7 Pines Meetings for the invitation, and to the other speakers (M. Devoret, J. Hartle, E. Heller, G. ‘t Hooft, D. Howard, M. Gutzwiller, M. Janssen, A. Leggett, R. Penrose, P. Stamp, and W. Zurek) for sharing their insights with him.

BIBLIOGRAPHY

- [Abraham and Marsden, 1985] R. Abraham and J. E. Marsden. *Foundations of Mechanics*, 2nd ed. Addison Wesley, Redwood City, 1985.
- [Accardi *et al.*, 1990] L. Accardi, A. Frigerio, and Y. Lu. The weak coupling limit as a quantum functional central limit. *Communications in Mathematical Physics*, 131: 537–570, 1990.
- [Adler, 2003] S. L. Adler. Why decoherence has not solved the measurement problem: A response to P. W. Anderson. *Studies in History and Philosophy of Modern Physics*, 34B: 135–142, 2003.
- [Agmon, 1982] S. Agmon. *Lectures on Exponential Decay of Solutions of Second-Order Elliptic Equations*. Princeton: Princeton University Press, 1982.
- [Albeverio and Høegh-Krohn, 1976] S. A. Albeverio and R. J. Høegh-Krohn. *Mathematical Theory of Feynman Path Integrals*. Berlin: Springer-Verlag, 1976.
- [Alfsen, 1970] E. M. Alfsen. *Compact Convex Sets and Boundary Integrals*. Berlin: Springer, 1970.

- [Ali *et al.*, 1995] S. T. Ali, J.-P. Antoine, J.-P. Gazeau, and U.A. Mueller. Coherent states and their generalizations: A mathematical overview. *Reviews in Mathematical Physics*, 7: 1013–1104, 1995.
- [Ali *et al.*, 2000] S. T. Ali, J.-P. Antoine, J.-P. Gazeau, and U.A. Mueller. *Coherent States, Wavelets and their Generalizations*. New York: Springer-Verlag, 2000.
- [Ali and Emch, 1986] S. T. Ali and G. G. Emch. Geometric quantization: modular reduction theory and coherent states. *Journal of Mathematical Physics*, 27: 2936–2943, 1986.
- [Ali and Englis, 2004] S. T. Ali and M. Englis. Quantization methods: a guide for physicists and analysts. 2004.
arXiv:math-ph/0405065
- [Alicki and Fannes, 2001] A. Alicki and M. Fannes. *Quantum Dynamical Systems*. Oxford: Oxford University Press, 2001.
- [Alicki and Lendi, 1987] A. Alicki and K. Lendi. *Quantum Dynamical Semigroups and Applications*. Berlin: Springer, 1987.
- [Amann, 1986] A. Amann. Observables in W^* -algebraic quantum mechanics. *Fortschritte der Physik*, 34: 167–215, 1986.
- [Amann, 1987] A. Amann. Broken symmetry and the generation of classical observables in large systems. *Helvetica Physica Acta*, 60: 384–393, 1987.
- [Amann and Primas, 1997] A. Amann and H. Primas. What is the referent of a non-pure quantum state? In S. Cohen, R. S., Horne, M. A., and J. Stachel (eds.), *Experimental Metaphysics: Quantum Mechanical Studies in Honor of Abner Shimony*. Dordrecht: Kluwer Academic Publishers, 1997.
- [Arai, 1995] T. Arai. Some extensions of the semiclassical limit $\hbar \rightarrow 0$ for Wigner functions on phase space. *Journal of Mathematical Physics*, 36: 622–630, 1995.
- [Araki, 1980] H. Araki. A remark on the Machida-Namiki theory of measurement. *Progress in Theoretical Physics*, 64: 719–730, 1980.
- [Araki, 1999] H. Araki. *Mathematical Theory of Quantum Fields*. New York: Oxford University Press, 1999.
- [Arnold, 1989] V. I. Arnold. *Mathematical Methods of Classical Mechanics*. Second edition. New York: Springer-Verlag, 1989.
- [Ashtekar and Schilling, 1999] A. Ashtekar and T. A. Schilling. Geometrical formulation of quantum mechanics. *On Einstein's Path (New York, 1996)*, pages 23–65. New York: Springer, 1999.
- [Atmanspacher *et al.*, 1989] H. Atmanspacher, A. Amann, and U. Müller-Herold (eds.). *On Quanta, Mind and Matter: Hans Primas in Context*. Dordrecht: Kluwer Academic Publishers, 1999.
- [Auletta, 2001] G. Auletta. *Foundations and Interpretation of Quantum Mechanics*. Singapore: World Scientific, 2001.
- [Bacciagaluppi, 1993] G. Bacciagaluppi. Separation theorems and Bell inequalities in algebraic quantum mechanics. In P. Busch, P. J. Lahti, and P. Mittelstaedt (eds.), *Proceedings of the Symposium on the Foundations of Modern Physics (Cologne, 1993)*, pages 29–37. Singapore: World Scientific, 1993.
- [Bacciagaluppi, 2004] G. Bacciagaluppi. The role of decoherence in quantum theory. In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, (Winter 2004 Edition), 2004.
Online only at <http://plato.stanford.edu/archives/win2004/entries/qm-decoherence/>.
- [Bach *et al.*, 1998] V. Bach, J. Fröhlich, and I. M. Sigal. Quantum electrodynamics of confined nonrelativistic particles. *Advances in Mathematics*, 137: 299–395, 1998.
- [Bach *et al.*, 1999] V. Bach, J. Fröhlich, and I. M. Sigal. Spectral analysis for systems of atoms and molecules coupled to the quantized radiation field. *Communications in Mathematical Physics*, 207: 249–290, 1999.
- [Baez, 1987] J. Baez. Bell's inequality for C^* -algebras. *Letters in Mathematical Physics*, 13: 135–136, 1987.
- [Bagarello and Morchio, 1992] F. Bagarello and G. Morchio. Dynamics of mean-field spin models from basic results in abstract differential equations. *Journal of Statistical Physics*, 66: 849–866, 1992.
- [Ballentine, 1970] L. E. Ballentine. The statistical interpretation of quantum mechanics. *Reviews of Modern Physics*, 42: 358–381, 1970.
- [Ballentine, 1986] L. E. Ballentine. Probability theory in quantum mechanics. *American Journal of Physics*, 54: 883–889, 1986.

- [Ballentine, 2002] L. E. Ballentine. Dynamics of quantum-classical differences for chaotic systems. *Physical Review*, A65, 062110-1–6, 2002.
- [Ballentine, 2003] L. E. Ballentine. The classical limit of quantum mechanics and its implications for the foundations of quantum mechanics. In A. Khrennikov (ed.), *Quantum Theory: Reconsideration of Foundations – 2*, pages 71–82. Växjö: Växjö University Press, 2003.
- [Ballentine et al., 1994] L. E. Ballentine, Y. Yang, and J. P. Zibin. Inadequacy of Ehrenfest’s theorem to characterize the classical regime. *Physical Review*, A50: 2854–2859, 1994.
- [Balian and Bloch, 1972] R. Balian and C. Bloch. Distribution of eigenfrequencies for the wave equation in a finite domain. III. Eigenfrequency density oscillations. *Annals of Physics*, 69: 76–160, 1972.
- [Balian and Bloch, 1974] R. Balian and C. Bloch. Solution of the Schrödinger equation in terms of classical paths. *Annals of Physics*, 85: 514–545, 1974.
- [Bambusi et al., 1999] D. Bambusi, S. Graffi, and T. Paul. Long time semiclassical approximation of quantum flows: a proof of the Ehrenfest time. *Asymptotic Analysis*, 21: 149–160, 1999.
- [Barrow-Green, 1997] J. Barrow-Green. *Poincaré and the Three Body Problem*. Providence, RI: (American Mathematical Society, 1997.
- [Barut and Račka, 1977] A. O. Barut and R. Račka. *Theory of Group Representations and Applications*. Warszawa: PWN, 1977.
- [Bassi and Ghirardi, 2000] A. Bassi and G. C. Ghirardi. Decoherent histories and realism. *Journal of Statistical Physics*, 98: 457–494. Reply by R. B. Griffiths. *ibid.*, 99: 1409–1425, 2000. Reply to this reply by A. Bassi and G. C. Ghirardi. *ibid.*, 99: 1427, 2000.
- [Bates and Weinstein, 1995] S. Bates and A. Weinstein. *Lectures on the Geometry of Quantization*. Berkeley Mathematics Lecture Notes, 8. University of California, Berkeley. Re-issued by the American Mathematical Society, 1995.
- [Batterman, 2002] R. W. Batterman. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford: Oxford University Press, 2002.
- [Batterman, 2005] R. W. Batterman. Critical phenomena and breaking drops: Infinite idealizations in physics. *Studies in History and Philosophy of Modern Physics*, 36: 225–244, 2005.
- [Baum et al., 1994] P. Baum, A. Connes, and N. Higson. Classifying space for proper actions and K-theory of group C^* -algebras. *Contemporary Mathematics*, 167: 241–291, 1994.
- [Bayen et al., 1978] F. Bayen, M. Flato, C. Fronsdal, A. Lichnerowicz, and D. Sternheimer. Deformation theory and quantization I, II. *Annals of Physics*, 110: 61–110, 111–151, 1978.
- [Bell, 1975] J. S. Bell. On wave packet reduction in the Coleman–Hepp model. *Helvetica Physica Acta*, 48: 93–98, 1975.
- [Bell, 1987] J. S. Bell. *Speakable and Unsayable in Quantum Mechanics*. Cambridge: Cambridge University Press, 1987.
- [Bell, 2001] J. S. Bell. *John S. Bell on the Foundations of Quantum Mechanics*. Singapore: World Scientific, 2001.
- [Beller, 1999] M. Beller. *Quantum Dialogue*. Chicago: University of Chicago Press, 1999.
- [Belot, 2005] G. Belot. The representation of time and change in mechanics. This volume, 2005.
- [Belot and Earman, 1997] G. Belot and J. Earman. Chaos out of order: quantum mechanics, the correspondence principle and chaos. *Studies in History and Philosophy of Modern Physics*, 28B: 147–182, 1997.
- [Beltrametti and Cassinelli, 1984] E. G. Beltrametti and G. Cassinelli. *The Logic of Quantum Mechanics*. Cambridge: Cambridge University Press, 1984.
- [Benatti, 1993] F. Benatti. *Deterministic Chaos in Infinite Quantum Systems*. Berlin: Springer-Verlag, 1993.
- [Bene and Dieks, 2002] G. Bene and D. Dieks. A perspectival version of the modal interpretation of quantum mechanics and the origin of macroscopic behavior. *Foundations of Physics*, 32: 645–671, 2002.
- [Berezin, 1974] F. A. Berezin. Quantization. *Mathematical USSR Izvestia*, 8: 1109–1163, 1974.
- [Berezin, 1975a] F. A. Berezin. Quantization in complex symmetric spaces. *Mathematical USSR Izvestia*, 9: 341–379, 1975a.
- [Berezin, 1975b] F. A. Berezin. General concept of quantization. *Communications in Mathematical Physics*, 40: 153–174, 1975b.
- [Berry, 1977a] M. V. Berry. Semi-classical mechanics in phase space: a study of Wigner’s function. *Philosophical Transactions of the Royal Society*, 287: 237–271, 1977a.

- [Berry, 1977b] M. V. Berry. Regular and irregular semi-classical wavefunctions. *Journal of Physics*, A10: 2083–2091, 1977b.
- [Berry et al., 1979] M. V. Berry, N. L. Balazs, M. Tabor, and A. Voros. Quantum maps. *Annals of Physics*, 122: 26–63, 1979.
- [Berry and Tabor, 1977] M. V. Berry and M. Tabor. Level clustering in the regular spectrum. *Proceedings of the Royal Society*, A356: 375–394, 1977.
- [Berry and Keating, 1999] M. V. Berry, and J. P. Keating. The Riemann zeros and eigenvalue asymptotics. *SIAM Review*, 41: 236–266, 1999.
- [Binz et al., 2004] E. Binz, R. Honegger, and A. Rieckers. Field-theoretic Weyl quantization as a strict and continuous deformation quantization. *Annales Henri Poincaré*, 5: 327–346, 2004.
- [Binz et al., 1988] E. Binz, J. Śniatycki, and H. Fischer. *The Geometry of Classical Fields*. Amsterdam: North–Holland, 1988.
- [Birkhoff and von Neumann, 1936] G. Birkhoff and J. von Neumann. The logic of quantum mechanics. *Annals of Mathematics*, (2) 37: 823–843, 1936.
- [Bitbol, 1996] M. Bitbol. *Schrödinger’s Philosophy of Quantum Mechanics*. Dordrecht: Kluwer Academic Publishers, 1996.
- [Benatti, 1992] M. Bitbol and O. Darrigol (eds.). *Erwin Schrödinger: Philosophy and the Birth of Quantum Mechanics*. Dordrecht: Kluwer Academic Publishers, 1992.
- [Blackadar, 1998] B. Blackadar. *K-Theory for Operator Algebras*. Second edition. Cambridge: Cambridge University Press, 1998.
- [Blair Bolles, 2004] E. Blair Bolles. *Einstein Defiant: Genius versus Genius in the Quantum Revolution*. Washington: Joseph Henry Press, 2004.
- [Blanchard, 1996] E. Blanchard. Deformations de C^* -algebras de Hopf. *Bulletin de la Société mathématique de France*, 124: 141–215, 1996.
- [Blanchard and Olkiewicz, 2003] Ph. Blanchard and R. Olkiewicz. Decoherence induced transition from quantum to classical dynamics. *Reviews in Mathematical Physics*, 15: 217–243, 2003.
- [Bohigas et al., 1984] O. Bohigas, M.-J. Giannoni, and C. Schmit. Characterization of chaotic quantum spectra and universality of level fluctuation laws. *Physical Review Letters*, 52: 1–4, 1984.
- [Blume-Kohout and Zurek, 2004] R. Blume-Kohout and W.H. Zurek. A simple example of “Quantum Darwinism”: Redundant information storage in many-spin environments. *Foundations of Physics*, to appear, 2004. [arXiv:quant-ph/0408147](https://arxiv.org/abs/quant-ph/0408147)
- [Blume-Kohout and Zurek, 2005] R. Blume-Kohout and W.H. Zurek. Quantum Darwinism: Entanglement, branches, and the emergent classicality of redundantly stored quantum information. *Physical Review A*, to appear, 2005. [arXiv:quant-ph/0505031](https://arxiv.org/abs/quant-ph/0505031)
- [Bogoliubov, 1958] N. N. Bogoliubov. On a new method in the theory of superconductivity. *Nuovo Cimento*, 7: 794–805, 1958.
- [Bohr, 1927] N. Bohr. The quantum postulate and the recent development of atomic theory. *Atti del Congresso Internazionale dei Fisici (Como, 1927)*. Reprinted in Bohr (1934), pages 52–91, 1927.
- [Bohr, 1934] N. Bohr. *Atomic Theory and the Description of Nature*. Cambridge: Cambridge University Press, 1934.
- [Bohr, 1935] N. Bohr. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 48: 696–702, 1935.
- [Bohr, 1937] N. Bohr. Causality and complementarity. *Philosophy of Science*, 4: 289–298, 1937.
- [Bohr, 1948] N. Bohr. On the notions of causality and complementarity. *Dialectica*, 2: 312–319, 1948.
- [Bohr, 1949] N. Bohr. Discussion with Einstein on epistemological problems in atomic physics. In P.A. Schlipp (ed.), *Albert Einstein: Philosopher-Scientist*, pages 201–241. La Salle: Open Court, 1949.
- [Bohr, 1958] N. Bohr. *Atomic Physics and Human Knowledge*. New York: Wiley, 1958.
- [Bohr, 1985] N. Bohr. In J. Kalckar (ed.), *Collected Works. Vol. 6: Foundations of Quantum Physics I (1926–1932)*. Amsterdam: North-Holland, 1985.
- [Bohr, 1996] N. Bohr. In J. Kalckar (ed.), *Collected Works. Vol. 7: Foundations of Quantum Physics II (1933–1958)*. Amsterdam: North-Holland, 1996.
- [Bona, 1980] P. Bona. A solvable model of particle detection in quantum theory. *Acta Facultatis Rerum Naturalium Universitatis Comenianae Physica*, XX: 65–94, 1980.

- [Bona, 1988] P. Bona. The dynamics of a class of mean-field theories. *Journal of Mathematical Physics*, 29: 2223–2235, 1988.
- [Bona, 1989] P. Bona. Equilibrium states of a class of mean-field theories. *Journal of Mathematical Physics*, 30: 2994–3007, 1989.
- [Bona, 2000] P. Bona. Extended quantum mechanics. *Acta Physica Slovaca*, 50: 1–198, 2000.
- [Bonechi and De Bièvre, 2000] F. Bonechi and S. De Bièvre. Exponential mixing and $\ln \hbar$ time scales in quantized hyperbolic maps on the torus. *Communications in Mathematical Physics*, 211: 659–686, 2000.
- [Bosse and Hartle, 2005] A. W. Bosse and J. B. Hartle. Representations of spacetime alternatives and their classical limits. 2005. [arXiv:quant-ph/0503182](https://arxiv.org/abs/quant-ph/0503182).
- [Brack and Bhaduri, 1997] M. Brack and R. K. Bhaduri. *Semiclassical Physics*. Boulder: Westview Press, 1997.
- [Bratteli and Robinson, 1987] O. Bratteli and D.W. Robinson. *Operator Algebras and Quantum Statistical Mechanics. Vol. I: C^* - and W^* -Algebras, Symmetry Groups, Decomposition of States*. 2nd Ed. Berlin: Springer, 1987.
- [Bratteli and Robinson, 1981] O. Bratteli and D.W. Robinson. *Operator Algebras and Quantum Statistical Mechanics. Vol. II: Equilibrium States, Models in Statistical Mechanics*. Berlin: Springer, 1981.
- [Brezger *et al.*, 2002] B. Brezger, L. Hackermüller, S. Uttenthaler, J. Petschinka, M. Arndt, and A. Zeilinger. Matter-Wave Interferometer for Large Molecules. *Physical Review Letters*, 88, 100404, 2002.
- [Breuer, 1994] T. Breuer. *Classical Observables, Measurement, and Quantum Mechanics*. Ph.D. Thesis, University of Cambridge, 1994.
- [Bröcker and Werner, 1995] T. Bröcker and R. F. Werner. Mixed states with positive Wigner functions. *Journal of Mathematical Physics*, 36: 62–75, 1995.
- [Brun and Hartle, 1999] T. A. Brun and J. B. Hartle. Classical dynamics of the quantum harmonic chain. *Physical Review*, D60, 123503–1–20, 1999.
- [Brush, 2002] S. G. Brush. Cautious revolutionaries: Maxwell, Planck, Hubble. *American Journal of Physics*, 70: 119–127, 2002.
- [Bub, 1988] J. Bub. How to Solve the Measurement Problem of Quantum Mechanics. *Foundations of Physics*, 18: 701–722, 1988.
- [Bub, 1999] J. Bub. *Interpreting the Quantum World*. Cambridge: Cambridge University Press, 1999.
- [Bub, 2004] J. Bub. Why the quantum? *Studies in History and Philosophy of Modern Physics*, 35B: 241–266, 2004.
- [Busch *et al.*, 1998] P. Busch, M. Grabowski, and P. J. Lahti. *Operational Quantum Physics*, 2nd corrected ed. Berlin: Springer, 1998.
- [Busch *et al.*, 1991] P. Busch, P. J. Lahti, and P. Mittelstaedt. *The Quantum Theory of Measurement*. Berlin: Springer, 1991.
- [Butterfield, 2002] J. Butterfield. Some Worlds of Quantum Theory. In R. Russell, J. Polkinghorne *et al.* (ed.). *Quantum Mechanics* (Scientific Perspectives on Divine Action vol 5), pages 111–140. Rome: Vatican Observatory Publications, 2, 2002. [arXiv:quant-ph/0105052](https://arxiv.org/abs/quant-ph/0105052); PITT-PHIL-SCI00000204.
- [Butterfield, 2005] J. Butterfield. On symmetry, conserved quantities and symplectic reduction in classical mechanics. *This volume*, 2005.
- [Büttner *et al.*, 2003] L. Büttner, J. Renn, and M. Schemmel. Exploring the limits of classical physics: Planck, Einstein, and the structure of a scientific revolution. *Studies in History and Philosophy of Modern Physics*, 34B: 37–60, 2003.
- [Camilleri, 2005] K. Camilleri. *Heisenberg and Quantum Mechanics: The Evolution of a Philosophy of Nature*. Ph.D. Thesis, University of Melbourne, 2005.
- [Cantoni, 1975] V. Cantoni. Generalized “transition probability”. *Communications in Mathematical Physics*, 44: 125–128, 1975.
- [Cantoni, 1977] V. Cantoni. The Riemannian structure on the states of quantum-like systems. *Communications in Mathematical Physics*, 56: 189–193, 1977.
- [Carson, 2000] C. Carson. Continuities and discontinuities in Planck’s *Akt der Verzweiflung*. *Annalen der Physik*, 9: 851–960, 2000.
- [Cassidy, 1992] D. C. Cassidy. *Uncertainty: the Life and Science of Werner Heisenberg*. New York: Freeman, 1992.

- [Castrigiano and Henrichs, 1980] D. P. L. Castrigiano and R. W. Henrichs. Systems of covariance and subrepresentations of induced representations. *Letters in Mathematical Physics*, 4: 169–175, 1980.
- [Cattaneo, 1979] U. Cattaneo. On Mackey’s imprimitivity theorem. *Commentari Mathematici Helvetici*, 54: 629–641, 1979.
- [Caves *et al.*, 2002] C. M. Caves, C. A. Fuchs, and R. Schack. Unknown quantum states: the quantum de Finetti representation. Quantum information theory. *Journal of Mathematical Physics*, 43: 4537–4559, 2002.
- [Charbonnel, 1986] A. M. Charbonnel. Localisation et développement asymptotique des éléments du spectre conjoint d’opérateurs pseudodifférentiels qui commutent. *Integral Equations Operator Theory*, 9: 502–536, 1986.
- [Charbonnel, 1988] A. M. Charbonnel. Comportement semi-classiques du spectre conjoint d’opérateurs pseudodifférentiels qui commutent. *Asymptotic Analysis*, 1: 227–261, 1988.
- [Charbonnel, 1992] A. M. Charbonnel. Comportement semi-classiques des systèmes ergodiques. *Annales de l’Institut Henri Poincaré - Physique Théorique*, 56: 187–214, 1992.
- [Chernoff, 1973] P. R. Chernoff. Essential self-adjointness of powers of generators of hyperbolic equations. *Journal of Functional Analysis*, 12: 401–414, 1973.
- [Chernoff, 1995] P. R. Chernoff. Irreducible representations of infinite dimensional transformation groups and Lie algebras I. *Journal of Functional Analysis*, 130: 255–282, 1995.
- [Chevalley, 1991] C. Chevalley. Introduction: Le dessin et la couleur. In E. Bauer and R. Omnès (eds.), *Niels Bohr: Physique Atomique et Connaissance Humaine*. (French translation of Bohr, 1958), pages 17–140. Paris: Gallimard, 1991.
- [Chevalley, 1999] C. Chevalley. Why do we find Bohr obscure? In D. Greenberger, W. L. Reiter, and A. Zeilinger (eds.), *Epistemological and Experimental Perspectives on Quantum Physics*, pages 59–74. Dordrecht: Kluwer Academic Publishers, 1999.
- [Chiorescu *et al.*, 2003] I. Chiorescu, Y. Nakamura, C. J. P. M. Harmans, and J. E. Mooij. Coherent quantum dynamics of a superconducting flux qubit. *Science*, 299, Issue 5614: 1869–1871, 2003.
- [Cirelli *et al.*, 1983] R. Cirelli, P. Lanzavecchia, and A. Mania. Normal pure states of the von Neumann algebra of bounded operators as a Kähler manifold. *Journal of Physics*, A16: 3829–3835, 1983.
- [Cirelli *et al.*, 1990] R. Cirelli, A. Maniá, and L. Pizzocchero. Quantum mechanics as an infinite-dimensional Hamiltonian system with uncertainty structure. I, II. *Journal of Mathematical Physics*, 31: 2891–2897, 2898–2903, 1990.
- [Colin de Verdière, 1973] Y. Colin de Verdière. Spectre du laplacien et longueurs des géodésiques périodiques. I, II. *Compositio Mathematica*, 27: 83–106, 159–184, 1973.
- [Colin de Verdière, 1977] Y. Colin de Verdière. Quasi-modes sur les variétés Riemanniennes. *Inventiones Mathematicae* 43, 15–52, 1977.
- [Colin de Verdière, 1985] Y. Colin de Verdière. Ergodicité et fonctions propres du Laplacien. *Communications in Mathematical Physics*, 102: 497–502, 1985.
- [Colin de Verdière, 1998] Y. Colin de Verdière. Une introduction à la mécanique semi-classique. *l’Enseignement Mathématique*, 44(2): 23–51, 1998.
- [Combesure, 1992] M. Combesure. The squeezed state approach of the semiclassical limit of the time-dependent Schrödinger equation. *Journal of Mathematical Physics*, 33: 3870–3880, 1992.
- [Combesure *et al.*, 1999] M. Combesure, J. Ralston, and D. Robert. A proof of the Gutzwiller semiclassical trace formula using coherent states decomposition. *Communications in Mathematical Physics*, 202: 463–480, 1999.
- [Combesure and Robert, 1997] M. Combesure and D. Robert. Semiclassical spreading of quantum wave packets and applications near unstable fixed points of the classical flow. *Asymptotic Analysis*, 14: 377–404, 1997.
- [Corwin and Greenleaf, 1989] L. Corwin and F. P. Greenleaf. *Representations of Nilpotent Lie Groups and Their Applications*, Part I. Cambridge: Cambridge University Press, 1989.
- [Cucchietti, 2004] F. M. Cucchietti. *The Loschmidt Echo in Classically Chaotic Systems: Quantum Chaos, Irreversibility and Decoherence*. Ph.D. Thesis, Universidad Nacional de Córdoba, 2004. [arXiv:quant-ph/0410121](https://arxiv.org/abs/quant-ph/0410121)
- [Cushing, 1994] J. T. Cushing. *Quantum Mechanics: Historical Contingency and the Copenhagen Hegemony*. Chicago: University of Chicago Press, 1994.

- [Cvitanovic *et al.*, 2005] P. Cvitanovic *et al.*. *Classical and Quantum Chaos*, 2005. <http://ChaosBook.org>.
- [Cycon *et al.*, 1987] H. L. Cycon, R. G. Froese, W. Kirsch, and B. Simon. *Schrödinger Operators with Application to Quantum Mechanics and Global Geometry*. Berlin: Springer-Verlag, 1987.
- [Darrigol, 1992] O. Darrigol. *From c-Numbers to q-Numbers*. Berkeley: University of California Press, 1992.
- [Darrigol, 2001] O. Darrigol. The Historians' Disagreements over the Meaning of Planck's Quantum. *Centaurus*, 43: 219–239, 2001.
- [Davidson, 2001] D. Davidson. *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press, 2001.
- [Davies, 1976] E. B. Davies. *Quantum Theory of Open Systems*. London: Academic Press, 1976.
- [De Bièvre, 1992] S. De Bièvre. Oscillator eigenstates concentrated on classical trajectories. *Journal of Physics*, A25: 3399–3418, 1992.
- [De Bièvre, 2001] S. De Bièvre. Quantum chaos: a brief first visit. *Contemporary Mathematics*, 289: 161–218, 2001.
- [De Bièvre, 2003] S. De Bièvre. Local states of free bose fields. Lectures given at the Summer School on Large Coulomb Systems, Nordfjordeid, 2003.
- [De Bièvre *et al.*, 1993] S. De Bièvre, M. Irac-Astaud, and J. C. Houard. Wave packets localised on closed classical trajectories. In W. F. Ames, E. M. Harrell, and J. V. Herod (eds.), *Differential Equations and Applications in Mathematical Physics*, pages 25–33. New York: Academic Press, 1993.
- [De Muynck, 2002] W. M. De Muynck. *Foundations of Quantum Mechanics: an Empiricist Approach*. Dordrecht: Kluwer Academic Publishers, 2002.
- [Devoret *et al.*, 2004] M. H. Devoret, A. Wallraff, and J. M. Martinis. Superconducting Qubits: A Short Review. 2004. [arXiv:cond-mat/0411174](https://arxiv.org/abs/cond-mat/0411174)
- [Diacu and Holmes, 1996] F. Diacu and P. Holmes. *Celestial Encounters. The Origins of Chaos and Stability*. Princeton: Princeton University Press, 1996.
- [Dickson, 2005] M. Dickson. Non-relativistic quantum mechanics. *This Volume*, 2005.
- [Dieks, 1989a] D. Dieks. Quantum mechanics without the projection postulate and its realistic interpretation. *Foundations of Physics*, 19: 1397–1423, 1989a.
- [Dieks, 1989b] D. Dieks. Resolution of the measurement problem through decoherence of the quantum state. *Physics Letters*, 142A: 439–446, 1989b.
- [Dimassi and Sjöstrand, 1999] M. Dimassi and J. Sjöstrand. *Spectral Asymptotics in the Semi-Classical Limit*. Cambridge: Cambridge University Press, 1999.
- [Dirac, 1926] P. A. M. Dirac. The fundamental equations of quantum mechanics. *Proceedings of the Royal Society*, A109: 642–653, 1926.
- [Dirac, 1930] P. A. M. Dirac. *The Principles of Quantum Mechanics*. Oxford: Clarendon Press, 1930.
- [Dirac, 1964] P. A. M. Dirac. *Lectures on Quantum Mechanics*. New York: Belfer School of Science, Yeshiva University, 1964.
- [Dixmier, 1977] J. Dixmier. *C*-Algebras*. Amsterdam: North-Holland, 1977.
- [Doebner and Tolar, 1975] H. D. Doebner and J. Tolar. Quantum mechanics on homogeneous spaces. *Journal of Mathematical Physics*, 16: 975–984, 1975.
- [Dowker and Kent, 1996] F. Dowker and A. Kent. On the consistent histories Approach to Quantum Mechanics. *Journal of Statistical Physics*, 82: 1575–1646, 1996.
- [Dubin *et al.*, 2000] D. A. Dubin, M. A. Hennings, and T. B. Smith. *Mathematical Aspects of Weyl Quantization and Phase*. Singapore: World Scientific, 2000.
- [Duclos and Høegre, 1993] P. Duclos and H. Høegre. On the semiclassical localization of the quantum probability. *Journal of Mathematical Physics*, 34: 1681–1691, 1993.
- [Duffield, 1990] N. G. Duffield. Classical and thermodynamic limits for generalized quantum spin systems. *Communications in Mathematical Physics*, 127: 27–39, 1990.
- [Duffield and Werner, 1992a] N. G. Duffield and R. F. Werner. Classical Hamiltonian dynamics for quantum Hamiltonian mean-field limits. In A. Truman and I. M. Davies (eds.), *Stochastics and Quantum Mechanics: Swansea, Summer 1990*, pages 115–129. Singapore: World Scientific, 1992a.
- [Duffield and Werner, 1992b] N. G. Duffield and R. F. Werner. On mean-field dynamical semi-groups on C^* -algebras. *Reviews in Mathematical Physics*, 4: 383–424, 1992b.
- [Duffield and Werner, 1992c] N. G. Duffield and R. F. Werner. Local dynamics of mean-field quantum systems. *Helvetica Physica Acta*, 65: 1016–1054, 1992c.

- [Duffield *et al.*, 1992] N. G. Duffield, H. Roos, and R. F. Werner. Macroscopic limiting dynamics of a class of inhomogeneous mean field quantum systems. *Annales de l'Institut Henri Poincaré - Physique Théorique*, 56: 143–186, 1992.
- [Duffner and Rieckers, 1988] E. Duffner and A. Rieckers. On the global quantum dynamics of multilattice systems with nonlinear classical effects. *Zeitschrift für Naturforschung*, A43: 521–532, 1988.
- [Duistermaat, 1974] J. J. Duistermaat. Oscillatory integrals, Lagrange immersions and unfolding of singularities. *Communications in Pure and Applied Mathematics*, 27: 207–281, 1974.
- [Duistermaat and Guillemin, 1975] J. J. Duistermaat and V. Guillemin. The spectrum of positive elliptic operators and periodic bicharacteristics. *Inventiones Mathematicae*, 29: 39–79, 1975.
- [Duistermaat, 1996] J. J. Duistermaat. *Fourier Integral Operators*. Original Lecture Notes from 1973. Basel: Birkhäuser, 1996.
- [Duval *et al.*, 1991] C. Duval, J. Elhadad, M. J. Gotay, J. Śniatycki, and G. M. Tuynman. Quantization and bosonic BRST theory. *Annals of Physics*, 206: 1–26, 1991.
- [Earman, 1986] J. Earman. *A Primer on Determinism*. Dordrecht: Reidel, 1986.
- [Earman, 2005] J. Earman. Aspects of determinism in modern physics. *This volume*, 2005.
- [Earman, to appear] J. Earman. Essential self-adjointness: implications for determinism and the classical-quantum correspondence. *Synthese*, to appear.
- [Echterhoff *et al.*, 2002] S. Echterhoff, S. Kaliszewski, J. Quigg, and I. Raeburn. A categorical approach to imprimitivity theorems for C^* -dynamical systems, 2002. [arXiv:math.0A/0205322](https://arxiv.org/abs/math/0602053)
- [Eddington, 1920] A. S. Eddington. *Space, Time, and Gravitation: An Outline of the General Relativity Theory*. Cambridge: Cambridge University Press, 1920.
- [Effros and Hahn, 1967] E. G. Effros and F. Hahn. Locally compact transformation groups and C^* -algebras. *Memoirs of the American Mathematical Society*, 75, 1967.
- [Ehrenfest, 1927] P. Ehrenfest. Bemerkung über die angenäherte Gültigkeit der klassischen Mechanik innerhalb der Quantenmechanik. *Zeitschrift für Physik*, 45, 455–457, 1927.
- [Einstein, 1905] A. Einstein. Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 17: 132–178, 1905.
- [Einstein, 1917] A. Einstein. Zum Quantensatz von Sommerfeld und Epstein. *Verhandlungen der deutschen Physikalischen Gesellschaft (2)*, 19: 82–92, 1917.
- [Einstein, 1949] A. Einstein. Remarks to the essays appearing in this collective volume. (Reply to criticisms). In P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist*, pages 663–688. La Salle: Open Court, 1949.
- [Elliott *et al.*, 1993] G. A. Elliott, T. Natsume, and R. Nest. The Heisenberg group and K -theory. *K-Theory*, 7: 409–428, 1993.
- [Emch and Knops, 1970] G. G. Emch and H. J. F. Knops. Pure thermodynamical phases as extremal KMS states. *Journal of Mathematical Physics*, 11: 3008–3018, 1970.
- [Emch, 1984] G. G. Emch. *Mathematical and conceptual foundations of 20th-century physics*. Amsterdam: North-Holland, 1984.
- [Emch and Liu, 2002] G. G. Emch and C. Liu. *The Logic of Thermostatistical Physics*. Berlin: Springer-Verlag, 2002.
- [Emch *et al.*, 1994] G. G. Emch, H. Narnhofer, W. Thirring, and G. Sewell. Anosov actions on noncommutative algebras. *Journal of Mathematical Physics*, 35: 5582–5599, 1994.
- [d’Espagnat, 1995] B. d’Espagnat. *Veiled Reality: An Analysis of Present-Day Quantum Mechanical Concepts*. Reading (MA): Addison-Wesley, 1995.
- [Esposito *et al.*, 2004] G. Esposito, G. Marmo, and G. Sudarshan. *From Classical to Quantum Mechanics: An Introduction to the Formalism, Foundations and Applications*. Cambridge: Cambridge University Press, 2004.
- [Enz, 2002] C. P. Enz. *No Time to be Brief: A Scientific Biography of Wolfgang Pauli*. Oxford: Oxford University Press, 2002.
- [Everett, 1957] H. Everett III. “Relative state” formulation of quantum mechanics. *Reviews in Modern Physics*, 29: 454–462, 1957.
- [Faye, 1991] J. Faye. *Niels Bohr: His Heritage and Legacy. An Anti-Realist View of Quantum Mechanics*. Dordrecht: Kluwer Academic Publishers, 1991.
- [Faye, 2002] J. Faye. Copenhagen interpretation of quantum mechanics. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2002 Edition), 2002. <http://plato.stanford.edu/archives/sum2002/entries/qm-copenhagen/>

- [Faye and Folse, 1994] J. Faye and H. Folse (eds.). *Niels Bohr and Contemporary Philosophy*. Dordrecht: Kluwer Academic Publishers, 1994.
- [Fell and Doran, 1988] J. M. G. Fell and R. S. Doran. *Representations of *-Algebras, Locally Compact Groups and Banach *-Algebraic Bundles*, Vol. 2. Boston: Academic Press, 1988.
- [Feyerabend, 1981] P. Feyerabend. Niels Bohr's world view. *Realism, Rationalism and Scientific Method: Philosophical Papers*, Vol. 1, pages 247–297. Cambridge: Cambridge University Press.
- [Fleming and Butterfield, 2000] G. Fleming and J. Butterfield. Strange positions. In J. Butterfield and C. Pagonis (eds.), *From Physics to Philosophy*, pages 108–165. Cambridge: Cambridge University Press, 2000.
- [Folse, 1985] H. J. Folse. *The Philosophy of Niels Bohr*. Amsterdam: North-Holland, 1985.
- [Ford, 1988] J. Ford. Quantum chaos. Is there any? In H. Bai-Lin (ed.), *Directions in Chaos*, Vol. 2, pages 128–147. Singapore: World Scientific, 1988.
- [Frasca, 2003] M. Frasca. General theorems on decoherence in the thermodynamic limit. *Physics Letters*, A308: 135–139, 2003.
- [Frasca, 2004] M. Frasca. Fully polarized states and decoherence. 2004. [arXiv:cond-mat/0403678](https://arxiv.org/abs/cond-mat/0403678)
- [Frigerio, 1974] A. Frigerio. Quasi-local observables and the problem of measurement in quantum mechanics. *Annales de l'Institut Henri Poincaré*, A3: 259–270, 1974.
- [Fröhlich et al., 2002] J. Fröhlich, T.-P. Tsai, and H.-T. Yau. On the point-particle (Newtonian) limit of the non-linear Hartree equation. *Communications in Mathematical Physics*, 225: 223–274, 2002.
- [Gallavotti, 1983] G. Gallavotti. *The Elements of Mechanics*. Berlin: Springer-Verlag, 1983.
- [Gallavotti et al., 2004] G. Gallavotti, F. Bonetto, and G. Gentile. *Aspects of Ergodic, Qualitative and Statistical Theory of Motion*. New York: Springer, 2004.
- [Gell-Mann and Hartle, 1990] M. Gell-Mann and J. B. Hartle. Quantum mechanics in the light of quantum cosmology. In W. H. Zurek (ed.) *Complexity, Entropy, and the Physics of Information*, pages 425–458. Reading, Addison-Wesley, 1990.
- [Gell-Mann and Hartle, 1993] M. Gell-Mann and J. B. Hartle. Classical equations for quantum systems. *Physical Review*, D47: 3345–3382, 1993.
- [Gérard and Leichtnam, 1993] P. Gérard and E. Leichtnam. Ergodic properties of eigenfunctions for the Dirichlet problem. *Duke Mathematical Journal*, 71: 559–607, 1993.
- [Gerisch et al., 1999] T. Gerisch, R. Münzner, and A. Rieckers. Global C^* -dynamics and its KMS states of weakly inhomogeneous bipolaronic superconductors. *Journal of Statistical Physics*, 97: 751–779, 1999.
- [Gerisch et al., 2003] T. Gerisch, R. Honegger, and A. Rieckers. Algebraic quantum theory of the Josephson microwave radiator. *Annales Henri Poincaré*, 4: 1051–1082, 2003.
- [Geyer et al., 1993] B. Geyer, H. Herwig, and H. Rechenberg (eds.). *Werner Heisenberg: Physiker und Philosoph*. Leipzig: Spektrum, 1993.
- [Giulini, 2003] D. Giulini. Superselection rules and symmetries. In E. Joos et al. (eds.), *Decoherence and the Appearance of a Classical World in Quantum Theory*, pages 259–316. Berlin: Springer, 2003.
- [Glimm and Jaffe, 1987] J. Glimm and A. Jaffe. *Quantum Physics. A Functional Integral Point of View*. New York: Springer-Verlag, 1987.
- [Gotay, 1986] M. J. Gotay. Constraints, reduction, and quantization. *Journal of Mathematical Physics*, 27: 2051–2066, 1986.
- [Gotay, 1999] M. J. Gotay. On the Groenewold-Van Hove problem for \mathbf{R}^{2n} . *Journal of Mathematical Physics*, 40: 2107–2116, 1999.
- [Gotay et al., 1996] M. J. Gotay, H. B. Grundling, and G. M. Tuynman. Obstruction results in quantization theory. *Journal of Nonlinear Science*, 6: 469–498, 1996.
- [Gotay et al., 1978] M. J. Gotay, J. M. Nester, and G. Hinds. Presymplectic manifolds and the Dirac–Bergmann theory of constraints. *Journal of Mathematical Physics*, 19: 2388–2399, 1978.
- [Götsch, 1992] J. Götsch. *Erwin Schrödinger's World View: The Dynamics of Knowledge and Reality*. Dordrecht: Kluwer Academic Publishers, 1992.
- [Govaerts, 1991] J. Govaerts. *Hamiltonian Quantization and Constrained Dynamics*. Leuven: Leuven University Press, 1991.
- [Gracia-Bondía et al., 2001] J. M. Gracia-Bondía, J. C. Várilly, and H. Figueroa. *Elements of Noncommutative Geometry*. Boston: Birkhäuser, 2001.

- [Griesemer *et al.*, 2001] M. Griesemer, E. H. Lieb, and M. Loss. Ground states in non-relativistic quantum electrodynamics. *Inventiones Mathematicae*, 145: 557–595, 2001.
- [Griffiths, 1984] R. B. Griffiths. Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics*, 36: 219–272, 1984.
- [Griffiths, 2002] R. B. Griffiths. *Consistent Quantum Theory*. Cambridge: Cambridge University Press, 2002.
- [Grigis and Sjöstrand, 1994] A. Grigis and J. Sjöstrand. *Microlocal Analysis for Differential Operators*. Cambridge: Cambridge University Press, 1994.
- [Groenewold, 1946] H. J. Groenewold. On the principles of elementary quantum mechanics. *Physica*, 12: 405–460, 1946.
- [Groenewold, 1952] H. J. Groenewold. Information in quantum measurements. *Proceedings Koninklijke Nederlandse Akademie van Wetenschappen*, B55: 219–227, 1952.
- [Gühr *et al.*, 1998] T. Guhr, H. Müller-Groeling, and H. Weidenmüller. Random matrix theories in quantum physics: common concepts. *Physics Reports*, 299: 189–425, 1998.
- [Guillemin *et al.*, 2002] V. Guillemin, V. Ginzburg, and Y. Karshon. *Moment Maps, Cobordisms, and Hamiltonian Group Actions*. Providence (RI): American Mathematical Society, 2002.
- [Guillemin and Sternberg, 1977] V. Guillemin and S. Sternberg. *Geometric Asymptotics*. Providence (RI): American Mathematical Society, 1977.
- [Guillemin and Sternberg, 1990] V. Guillemin and S. Sternberg. *Variations on a Theme by Kepler*. Providence (RI): American Mathematical Society, 1990.
- [Guillemin and Uribe, 1989] V. Guillemin and A. Uribe. Circular symmetry and the trace formula. *Inventiones Mathematicae*, 96: 385–423, 1989.
- [Gustafson and Sigal, 2003] S. J. Gustafson and I. M. Sigal. *Mathematical concepts of quantum mechanics*. Berlin: Springer, 2003.
- [Gutzwiller, 1971] M. C. Gutzwiller. Periodic orbits and classical quantization conditions. *Journal of Mathematical Physics*, 12: 343–358, 1971.
- [Gutzwiller, 1990] M. C. Gutzwiller. *Chaos in Classical and Quantum Mechanics*. New York: Springer-Verlag, 1990.
- [Gutzwiller, 1992] M. C. Gutzwiller. Quantum chaos. *Scientific American*, 266: 78–84, 1992.
- [Gutzwiller, 1998] M. C. Gutzwiller. Resource letter ICQM-1: The interplay between classical and quantum mechanics. *American Journal of Physics*, 66: 304–324, 1998.
- [Haag, 1962] R. Haag. The mathematical structure of the Bardeen–Cooper–Schrieffer model. *Nuovo Cimento*, 25: 287–298, 1962.
- [Haag *et al.*, 1970] R. Haag, R. Kadison, and D. Kastler. Nets of C^* -algebras and classification of states. *Communications in Mathematical Physics*, 16: 81–104, 1970.
- [Haag, 1992] R. Haag. *Local Quantum Physics: Fields, Particles, Algebras*. Heidelberg: Springer-Verlag, 1992.
- [Haake, 2001] F. Haake. *Quantum Signatures of Chaos*. Second Edition. New York: Springer-Verlag, 2001.
- [Hagedorn, 1998] G. A. Hagedorn. Raising and lowering operators for semiclassical wave packets. *Annals of Physics*, 269: 77–104, 1998.
- [Hagedorn and Joye, 1999] G. A. Hagedorn and A. Joye. Semiclassical dynamics with exponentially small error estimates. *Communications in Mathematical Physics*, 207: 439–465, 1999.
- [Hagedorn, 2000] G. A. Hagedorn and A. Joye. Exponentially accurate semiclassical dynamics: propagation, localization, Ehrenfest times, scattering, and more general states. *Annales Henri Poincaré*, 1: 837–883, 2000.
- [Halliwell, 1998] J. J. Halliwell. Decoherent histories and hydrodynamic equations. *Physical Review*, D58, 105015-1–12, 1998.
- [Halliwell, 2000] J. J. Halliwell. The emergence of hydrodynamic equations from quantum theory: a decoherent histories analysis. *International Journal of Theoretical Physics*, 39: 1767–1777, 2000.
- [Halliwell, 2004] J. J. Halliwell. Some recent developments in the decoherent histories approach to quantum theory. *Lecture Notes in Physics*, 633: 63–83, 2004.
- [Halliwell, 2005] J. J. Halliwell. How the quantum Universe becomes classical. 2005. [arXiv:quant-ph/0501119](https://arxiv.org/abs/quant-ph/0501119)
- [Halvorson, 2004] H. Halvorson. Complementarity of representations in quantum mechanics. *Studies in History and Philosophy of Modern Physics*, B35: 45–56, 2004.
- [Halvorson, 2005] H. Halvorson. Algebraic quantum field theory. This volume, 2005.

- [Halvorson and Clifton, 1999] H. Halvorson and R. Clifton. Maximal beable subalgebras of quantum-mechanical observables. *International Journal of Theoretical Physics*, 38: 2441–2484, 1999.
- [Halvorson and Clifton, 2002] H. Halvorson and R. Clifton. Reconsidering Bohr’s reply to EPR. In T. Placek and J. Butterfield (eds.), *Non-locality and Modality*, pages 3–18. Dordrecht: Kluwer Academic Publishers, 2002.
- [Hannabuss, 1984] K. C. Hannabuss. Dilations of a quantum measurement. *Helvetica Physica Acta*, 57: 610–620, 1984.
- [Harrison and Wan, 1997] F. E. Harrison and K. K. Wan. Macroscopic quantum systems as measuring devices: dc SQUIDs and superselection rules. *Journal of Physics*, A30: 4731–4755, 1997.
- [Hartle, 1995] J. B. Hartle. Spacetime quantum mechanics and the quantum mechanics of spacetime. *Gravitation et Quantifications (Les Houches, 1992)*, page 285–480. Amsterdam: North-Holland, 1995.
- [Hartle, 2005] J. B. Hartle. What connects different interpretations of quantum mechanics? In A. Elitzur, S. Dolev, and N. Kolenda (eds.), *Quo Vadis Quantum Mechanics*, pages 73–82. Heidelberg: Springer-Verlag, 2005. [arXiv:quant-ph/0305089](https://arxiv.org/abs/quant-ph/0305089)
- [Hawkins, 2005] E. Hawkins. Quantization of multiply connected manifolds. *Communications in Mathematical Physics*, 255: 513–575, 2005.
- [Heath and Sudderth, 1976] D. Heath and W. Sudderth. De Finetti’s theorem on exchangeable variables. *American Statistics*, 30: 188–189, 1976.
- [Heelan, 1965] P. Heelan. *Quantum Mechanics and Objectivity: A Study of the Physical Philosophy of Werner Heisenberg*. Den Haag: Martinus Nijhoff, 1965.
- [Heilbron, 2000] Heilbron, J. *The Dilemmas of an Upright Man: Max Planck as a Spokesman for German Science*. Second Edition. Los Angeles: University of California Press, 2000.
- [Heisenberg, 1925] W. Heisenberg. Über die quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen. *Zeitschrift für Physik*, 33: 879–893, 1925.
- [Heisenberg, 1927] W. Heisenberg. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Zeitschrift für Physik*, 43: 172–198, 1927.
- [Heisenberg, 1930] W. Heisenberg. *The Physical Principles of the Quantum Theory*. Chicago: University of Chicago Press, 1930.
- [Heisenberg, 1942] W. Heisenberg. Ordnung der Wirklichkeit. In Heisenberg [1984a], pages 217–306, 1942. Also available at <http://werner-heisenberg.unh.edu/Ordnung.htm>.
- [Heisenberg, 1958] W. Heisenberg. *Physics and Philosophy: The Revolution in Modern Science*. London: Allen and Unwin, 1958.
- [Heisenberg, 1969] W. Heisenberg. *Der Teil und das Ganze: Gespräche im Umkreis der Atomphysik*. München: Piper, 1969. English translation as Heisenberg [1971].
- [Heisenberg, 1971] W. Heisenberg. *Physics and Beyond*. New York: Harper and Row, 1971. Translation of Heisenberg [1969].
- [Heisenberg, 1984a] W. Heisenberg. *Gesammelte Werke. Series C: Philosophical and Popular Writings, Vol I: Physik und Erkenntnis 1927–1955*. W. Blum, H.-P. Dürr, and H. Rechenberg (eds.). München: Piper, 1984a.
- [Heisenberg, 1984b] W. Heisenberg. *Gesammelte Werke. Series C: Philosophical and Popular Writings, Vol II: Physik und Erkenntnis 1956–1968*. W. Blum, H.-P. Dürr, and H. Rechenberg (eds.). München: Piper, 1984b.
- [Heisenberg, 1985] W. Heisenberg. *Gesammelte Werke. Series C: Philosophical and Popular Writings, Vol III: Physik und Erkenntnis 1969–1976*. W. Blum, H.-P. Dürr, and H. Rechenberg (eds.). München: Piper, 1985.
- [Held, 1994] C. Held. The Meaning of complementarity. *Studies in History and Philosophy of Science*, 25: 871–893, 1994.
- [Helffer, 1988] B. Helffer. *Semi-classical Analysis for the Schrödinger Operator and Applications. Lecture Notes in Mathematics*, 1336. Berlin: Springer-Verlag, 1988.
- [Heller and Tomsovic, 1993] E. J. Heller and S. Tomsovic. Postmodern quantum mechanics. *Physics Today*, July: 38–46.
- [Hendry, 1984] J. Hendry. *The Creation of Quantum Mechanics and the Bohr-Pauli Dialogue*. Dordrecht: D. Reidel, 1984.
- [Henneaux and Teitelboim, 1992] M. Henneaux and C. Teitelboim. *Quantization of Gauge Systems*. Princeton: Princeton University Press, 1992.

- [Hepp, 1972] K. Hepp. Quantum theory of measurement and macroscopic observables. *Helvetica Physica Acta*, 45: 237–248, 1972.
- [Hepp, 1974] K. Hepp. The classical limit of quantum mechanical correlation functions. *Communications in Mathematical Physics*, 35: 265–277, 1974.
- [Hepp and Lieb, 1974] K. Hepp and E. Lieb. Phase transitions in reservoir driven open systems with applications to lasers and superconductors. *Helvetica Physica Acta*, 46: 573–602, 1974.
- [Higson, 1990] N. Higson. A primer on KK -theory. *Operator Theory: Operator Algebras and Applications. Proceedings Symposia in Pure Mathematical*, 51, Part 1, pages 239–283. Providence, RI: American Mathematical Society, 1990.
- [Hillery *et al.*, 1984] M. Hillery, R. F. O’Connell, M. O. Scully, and E. P. Wigner. Distribution functions in physics – Fundamentals. *Physics Reports*, 106: 121–167, 1984.
- [Hislop and Sigal, 1996] P. D. Hislop and I. M. Sigal. *Introduction to Spectral Theory. With Applications to Schrödinger Operators*. New York: Springer-Verlag, 1996.
- [Hoffmann, 1810] E. T. A. Hoffmann. *Musikalische Novellen und Aufsätze*. Leipzig: Insel-Bücherei, 1810.
- [Holevo, 1982] A. S. Holevo. *Probabilistic and Statistical Aspects of Quantum Theory*. Amsterdam: North-Holland Publishing Co, 1982.
- [Hogreve *et al.*, 1983] H. Hogreve, J. Potthoff, and R. Schrader. Classical limits for quantum particles in external Yang–Mills potentials. *Communications in Mathematical Physics*, 91: 573–598, 1983.
- [Honegger and Rieckers, 1994] R. Honegger and A. Rieckers. Quantized radiation states from the infinite Dicke model. *Publications of the Research Institute for Mathematical Sciences (Kyoto)*, 30: 111–138, 1994.
- [Honner, 1987] J. Honner. *The Description of Nature: Niels Bohr and the Philosophy of Quantum Physics*. Oxford: Oxford University Press, 1987.
- [Hooker, 1972] C. A. Hooker. The nature of quantum mechanical reality: Einstein versus Bohr. In J. Colodny (ed.). *Paradigms and Paradoxes: The Philosophical Challenges of the Quantum Domain*, pages 67–302. Pittsburgh: University of Pittsburgh Press, 1972.
- [Hörmander, 1965] L. Hörmander. Pseudo-differential operators. *Communications in Pure Applied Mathematical*, 18: 501–517, 1965.
- [Hörmander, 1979] L. Hörmander. The Weyl calculus of pseudo-differential operators. *Communications in Pure Applied Mathematical*, 32: 359–443, 1979.
- [Hörmander, 1985a] L. Hörmander. *The Analysis of Linear Partial Differential Operators, Vol. III*. Berlin: Springer-Verlag, 1985a.
- [Hörmander, 1985b] L. Hörmander. *The Analysis of Linear Partial Differential Operators, Vol. IV*. Berlin: Springer-Verlag, 1985b.
- [Horowitz and Marolf, 1995] G. T. Horowitz and D. Marolf. Quantum probes of spacetime singularities. *Physical Review*, D52: 5670–5675, 1995.
- [Hörz, 1968] H. Hörz. *Werner Heisenberg und die Philosophie*. Berlin: VEB Deutscher Verlag der Wissenschaften, 1968.
- [Howard, 1990] D. Howard. ‘Nicht sein kann was nicht sein darf’, or the Prehistory of EPR, 1909–1935: Einstein’s early worries about the quantum mechanics of composite systems. In A. I. Miller (ed.), *Sixty-Two Years of Uncertainty*, pages 61–111. New York: Plenum, 1990.
- [Howard, 1994] D. Howard. What makes a classical concept classical? Towards a reconstruction of Niels Bohr’s philosophy of physics. In J. Faye and H. Folse (eds.), *Niels Bohr and Contemporary Philosophy*, pages 201–229. Dordrecht: Kluwer Academic Publishers, 1994.
- [Howard, 2004] D. Howard. Who Invented the Copenhagen Interpretation? *Philosophy of Science*, 71: 669–682, 2004.
- [Howe, 1980] R. Howe. Quantum mechanics and partial differential equations. *Journal of Functional Analysis*, 38: 188–254, 1980.
- [Hudson, 1974] R. L. Hudson. When is the Wigner quasi-probability density non-negative? *Reports of Mathematical Physics*, 6: 249–252, 1974.
- [Hudson and Moody, 1975/76] R. L. Hudson and G. R. Moody. Locally normal symmetric states and an analogue of de Finetti’s theorem. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 33: 343–351, 1975/76.
- [Hunziker and Sigal, 2000] W. Hunziker and I. M. Sigal. The quantum N -body problem. *Journal of Mathematical Physics*, 41: 3448–3510, 2000.
- [Husimi, 1940] K. Husimi. Some formal properties of the density matrix. *Progress of the Physical and Mathematical Society of Japan*, 22: 264–314, 1940.

- [Isham, 1984] C. J. Isham. Topological and global aspects of quantum theory. *Relativity, Groups and Topology, II (Les Houches, 1983)*, pages 1059–1290. Amsterdam: North-Holland, 1984.
- [Isham, 1994] C. J. Isham. Quantum logic and the histories approach to quantum theory. *Journal of Mathematical Physics*, 35: 2157–2185, 1994.
- [Isham, 1997] C. J. Isham. Topos theory and consistent histories: the internal logic of the set of all consistent sets. *International Journal of Theoretical Physics*, 36: 785–814, 1997.
- [Isham and Butterfield, 2000] C. J. Isham and J. Butterfield. Some possible roles for topos theory in quantum theory and quantum gravity. *Foundations of Physics*, 30: 1707–1735, 2000.
- [Isham and Linden, 1994] C. J. Isham and N. Linden. Quantum temporal logic and decoherence functionals in the histories approach to generalized quantum theory. *Journal of Mathematical Physics*, 35: 5452–5476, 1994.
- [Isham and Linden, 1995] C. J. Isham and N. Linden. Continuous histories and the history group in generalized quantum theory. *Journal of Mathematical Physics*, 36: 5392–5408, 1995.
- [Isham et al., 1994] C. J. Isham, N. Linden, and S. Schreckenberg. The classification of decoherence functionals: an analog of Gleason’s theorem. *Journal of Mathematical Physics*, 35: 6360–6370, 1994.
- [Ivrii, 1998] V. Ivrii. *Microlocal Analysis and Precise Spectral Asymptotics*. New York: Springer-Verlag, 1998.
- [Jalabert and Pastawski, 2001] R. O. Jalabert and H. M. Pastawski. Environment-Independent Decoherence Rate in Classically Chaotic Systems. *Physical Review Letters*, 86: 2490–2493, 2001.
- [Jammer, 1966] M. Jammer. *The Conceptual Development of Quantum Mechanics*. New York: McGraw-Hill, 1966.
- [Jammer, 1974] M. Jammer. *The Philosophy of Quantum Mechanics*. New York: Wiley, 1974.
- [Janssens, 2005] B. Janssens. *Quantum Measurement: A Coherent Description*. M.Sc. Thesis, Radboud Universiteit Nijmegen, 2005. [arXiv:quant-ph/0503009](https://arxiv.org/abs/quant-ph/0503009)
- [Jauch, 1968] J. M. Jauch. *Foundations of Quantum Mechanics*. Reading (MA): Addison-Wesley, 1968.
- [Joos and Zeh, 1985] E. Joos and H. D. Zeh. The emergence of classical properties through interaction with the environment. *Zeitschrift für Physik*, B59: 223–243, 1985.
- [Joos et al., 2003] E. Joos, H. D. Zeh, C. Kiefer, D. Giulini, J. Kupsch, and I.-O. Stamatescu. *Decoherence and the Appearance of a Classical World in Quantum Theory*. Berlin: Springer-Verlag, 2003.
- [Jørgensen and Moore, 1984] P. E. T. Jørgensen and R. T. Moore. *Operator Commutation Relations*. Dordrecht: Reidel, 1984.
- [Kaplan and Heller, 1998a] L. Kaplan and E. J. Heller. Linear and nonlinear theory of eigenfunction scars. *Annals of Physics*, 264: 171–206, 1998a.
- [Kaplan and Heller, 1998b] L. Kaplan and E. J. Heller. Weak quantum ergodicity. *Physica D*, 121: 1–18, 1998b.
- [Kaplan, 1999] L. Kaplan. Scars in quantum chaotic wavefunctions. *Nonlinearity*, 12: R1–R40, 1999.
- [Katok and Hasselblatt, 1995] A. Katok and B. Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge: Cambridge University Press, 1995.
- [Kadison and Ringrose, 1983] R. V. Kadison and J. R. Ringrose. *Fundamentals of the theory of operator algebras. Vol. 1: Elementary Theory*. New York: Academic Press, 1983.
- [Kadison and Ringrose, 1986] R. V. Kadison and J. R. Ringrose. *Fundamentals of the theory of operator algebras. Vol. 2: Advanced Theory*. New York: Academic Press, 1986.
- [Karasev, 1989] M. V. Karasev. The Maslov quantization conditions in higher cohomology and analogs of notions developed in Lie theory for canonical fibre bundles of symplectic manifolds. I, II. *Selecta Mathematica Formerly Sovietica*, 8: 213–234, 235–258, 1989.
- [Kent, 1990] A. Kent. Against Many-Worlds Interpretations. *International Journal of Modern Physics*, A5, 1745, 1990.
- [Kent, 1997] A. Kent. Consistent sets yield contrary inferences in quantum theory. *Physical Review Letters*, 78: 2874–2877, 1997. Reply by R. B. Griffiths and J. B. Hartle (1998). *ibid.* 81, 1981. Reply to this reply by A. Kent (1998). *ibid.* 81, 1982.
- [Kent, 1998] A. Kent. Quantum histories. *Physica Scripta*, T76: 78–84, 1998.
- [Kent, 2000] A. Kent. Night thoughts of a quantum physicist. *Philosophical Transactions of the Royal Society of London*, A 358: 75–88, 2000.

- [Kiefer, 2003] C. Kiefer. Consistent histories and decoherence. In E. Joos *et al.* (eds.), *Decoherence and the Appearance of a Classical World in Quantum Theory*, pages 227–258. Berlin: Springer-Verlag, 2003.
- [Kirchberg and Wassermann, 1995] E. Kirchberg and S. Wassermann. Operations on continuous bundles of C^* -algebras. *Mathematische Annalen*, 303: 677–697, 1995.
- [Kirillov, 1990] A. A. Kirillov. Geometric quantization. In V. I. Arnold and S.P. Novikov (eds.), *Dynamical Systems IV*, pages 137–172. Berlin: Springer-Verlag, 1990.
- [Kirillov, 2004] A. A. Kirillov. *Lectures on the Orbit Method*. Providence, RI: American Mathematical Society, 2004.
- [Klauder and Skagerstam, 1985] J. R. Klauder and B.-S. Skagerstam (eds.). *Coherent States*. Singapore: World Scientific, 1985.
- [Klingenberg, 1982] W. Klingenberg. *Riemannian Geometry*. de Gruyter, Berlin, 1982.
- [Kohn and Nirenberg, 1965] J. J. Kohn and L. Nirenberg. An algebra of pseudo-differential operators. *Communications in Pure and Applied Mathematics*, 18: 269–305, 1965.
- [Koopman, 1931] B. O. Koopman. Hamiltonian systems and transformations in Hilbert space. *Proceedings of the National Academy of Sciences*, 18: 315–318, 1931.
- [Kostant, 1970] B. Kostant. Quantization and unitary representations. *Lecture Notes in Mathematics*, 170: 87–208, 1970.
- [Krishnaprasad and Marsden, 1987] P. S. Krishnaprasad and J. E. Marsden. Hamiltonian structures and stability for rigid bodies with flexible attachments. *Archive of Rational Mechanics and Analysis*, 98: 71–93, 1987.
- [Kuhn, 1978] T. S. Kuhn. *Black-body Theory and the Quantum Discontinuity: 1894–1912*. New York: Oxford University Press, 1978.
- [Kümmerer, 2002] B. Kümmerer. Quantum Markov processes. A. Buchleitner and K. Hornberger (eds.), *Coherent Evolution in Noisy Environment*. Lecture Notes in Physics, Vol. 611, pages, 139–198. Berlin: Springer-Verlag, 2002.
- [Lahti and Mittelstaedt, 1987] P. Lahti and P. Mittelstaedt (eds.). *The Copenhagen Interpretation 60 Years After the Como Lecture*. Singapore: World Scientific, 1987.
- [Landau and Lifshitz, 1977] L. D. Landau and E. M. Lifshitz. *Quantum Mechanics: Non-relativistic Theory*. 3d Ed. Oxford: Pergamon Press, 1977.
- [Landsman, 1990a] N. P. Landsman. Quantization and superselection sectors I. Transformation group C^* -algebras. *Reviews in Mathematical Physics*, 2: 45–72, 1990a.
- [Landsman, 1990b] N. P. Landsman. Quantization and superselection sectors II. Dirac Monopole and Aharonov–Bohm effect. *Reviews in Mathematical Physics*, 2: 73–104, 1990b.
- [Landsman, 1991] N. P. Landsman. Algebraic theory of superselection sectors and the measurement problem in quantum mechanics. *International Journal of Modern Physics*, A30: 5349–5371, 1991.
- [Landsman, 1992] N. P. Landsman. Induced representations, gauge fields, and quantization on homogeneous spaces. *Reviews in Mathematical Physics*, 4: 503–528, 1992.
- [Landsman, 1993] N. P. Landsman. Deformations of algebras of observables and the classical limit of quantum mechanics. *Reviews in Mathematical Physics*, 5: 775–806, 1993.
- [Landsman, 1995] N. P. Landsman. Observation and superselection in quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 26B: 45–73, 1995.
- [Landsman, 1997] N. P. Landsman. Poisson spaces with a transition probability. *Reviews in Mathematical Physics*, 9: 29–57, 1997.
- [Landsman, 1998] N. P. Landsman. *Mathematical Topics Between Classical and Quantum Mechanics*. New York: Springer-Verlag, 1998.
- [Landsman, 1999a] N. P. Landsman. Quantum mechanics on phase space. *Studies in History and Philosophy of Modern Physics*, 30B: 287–305, 1999a.
- [Landsman, 1999b] N. P. Landsman. Lie groupoid C^* -algebras and Weyl quantization. *Communications in Mathematical Physics*, 206: 367–381, 1999b.
- [Landsman, 2001] N. P. Landsman. Quantized reduction as a tensor product. In N. P. Landsman, M. J. Pflaum, and M. Schlichenmaier (eds.), *Quantization of Singular Symplectic Quotients*, pages 137–180. Basel: Birkhäuser, 2001.
- [Landsman, 2002] N. P. Landsman. Quantization as a functor. *Contemporary Mathematics*, 315: 9–24, 2002.

- [Landsman, 2005] N. P. Landsman. Functorial quantization and the Guillemin–Sternberg conjecture. S. T. Ali, G. G. Emch, A. Odziejewicz, M. Schlichenmaier, and S. L. Woronowicz (eds). *Twenty Years of Białowieża: A Mathematical Anthology*, pages 23–45. Singapore: World Scientific, 2005. [arXiv:math-ph/0307059](#)
- [Landsman, 2006a] N. P. Landsman. Lie Groupoids and Lie algebroids in physics and noncommutative geometry. *Journal of Geom. Physics*, 56: 24–54, 2006a.
- [Landsman, 2006b] N. P. Landsman. When champions meet: Rethinking the Bohr–Einstein debate. *Studies in History and Philosophy of Modern Physics*, 37:212–241, 2006b. [arXiv:quant-ph/0507220](#)
- [Landsman and Ramazan, 2001] N. P. Landsman and B. Ramazan. Quantization of Poisson algebras associated to Lie algebroids. *Contemporary Mathematics*, 282: 159–192, 2001.
- [Laurikainen, 1988] K. V. Laurikainen. *Beyond the Atom: The Philosophical Thought of Wolfgang Pauli*. Berlin: Springer-Verlag, 1988.
- [Lazutkin, 1993] V. F. Lazutkin. *KAM Theory and Semiclassical Approximations to Eigenfunctions*. Berlin: Springer-Verlag, 1993.
- [Leggett, 2002] A. J. Leggett. Testing the limits of quantum mechanics: motivation, state of play, prospects. *Journal of Physics: Condensed Matter*, 14: R415–R451, 2002.
- [Liboff, 1984] R. L. Liboff. The correspondence principle revisited. *Physics Today*, February: 50–55, 1984.
- [Lieb, 1973] E. H. Lieb. The classical limit of quantum spin systems. *Communications in Mathematical Physics*, 31: 327–340, 1973.
- [Littlejohn, 1986] R. G. Littlejohn. The semiclassical evolution of wave packets. *Physics Reports*, 138: 193–291, 1986.
- [Ludwig, 1985] G. Ludwig. *An Axiomatic Basis for Quantum Mechanics. Volume 1: Derivation of Hilbert Space Structure*. Berlin: Springer-Verlag, 1985.
- [Lugiewicz and Olkiewicz, 2002] P. Lugiewicz and R. Olkiewicz. Decoherence in infinite quantum systems. *Journal of Physics*, A35: 6695–6712, 2002.
- [Lugiewicz and Olkiewicz, 2003] P. Lugiewicz and R. Olkiewicz. Classical properties of infinite quantum open systems. *Communications in Mathematical Physics*, 239: 241–259, 2003.
- [Maassen, 2003] H. Maassen. Quantum probability applied to the damped harmonic oscillator. *Quantum Probability Communications*, Vol. XII (Grenoble, 1998), pages 23–58. River Edge, NJ: World Scientific Publishing, 2003.
- [Mackey, 1962] G. W. Mackey. *The Mathematical Foundations of Quantum Mechanics*. New York: Benjamin, 1962.
- [Mackey, 1968] G. W. Mackey. *Induced Representations of Groups and Quantum Mechanics*. New York: W. A. Benjamin; Turin: Editore Boringhieri, 1968.
- [Mackey, 1978] G. W. Mackey. *Unitary Group Representations in Physics, Probability, and Number Theory*. Reading, Mass.: Benjamin/Cummings Publishing Co, 1978.
- [Mackey, 1992] G. W. Mackey. *The Scope and History of Commutative and Noncommutative Harmonic Analysis*. Providence, RI: American Mathematical Society, 1992.
- [Majid, 1988] S. Majid. Hopf algebras for physics at the Planck scale. *Classical and Quantum Gravity*, 5: 1587–1606, 1988.
- [Majid, 1990] S. Majid. Physics for algebraists: noncommutative and noncocommutative Hopf algebras by a bicrossproduct construction. *Journal of Algebra*, 130: 17–64, 1990.
- [Marmo *et al.*, 2005] G. Marmo, G. Scolarici, A. Simoni, and F. Ventriglia. The quantum-classical transition: the fate of the complex structure. *International Journal of Geometric Methods in Physics*, 2: 127–145, 2005.
- [Marsden, 1992] J. E. Marsden. *Lectures on Mechanics*. Cambridge: Cambridge University Press, 1992.
- [Marsden and Ratiu, 1994] J. E. Marsden and T. S. Ratiu. *Introduction to Mechanics and Symmetry*. New York: Springer-Verlag, 1994.
- [Marsden *et al.*, 1984] J. E. Marsden, T. Ratiu, and A. Weinstein. Semidirect products and reduction in mechanics. *Transactions of the American Mathematical Society*, 281: 147–177, 1984.
- [Marshall *et al.*, 2003] W. Marshall, C. Simon, R. Penrose, and D. Bouwmeester. Towards quantum superpositions of a mirror. *Physical Review Letters*, 91, 130401-1–4, 2003.
- [Martinez, 2002] A. Martinez. *An Introduction to Semiclassical and Microlocal Analysis*. New York: Springer-Verlag, 2002.

- [Maslov and Fedoriuk, 1981] V. P. Maslov and M. V. Fedoriuk. *Semi-Classical Approximation in Quantum Mechanics*. Dordrecht: Reidel.
- [McCormmach, 1982] R. McCormmach. *Night Thoughts of a Classical Physicist*. Cambridge (MA): Harvard University Press, 1982.
- [Mehra and Rechenberg, 1982a] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 1: The Quantum Theory of Planck, Einstein, Bohr, and Sommerfeld: Its Foundation and the Rise of Its Difficulties*. New York: Springer-Verlag, 1982a.
- [Mehra and Rechenberg, 1982b] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 2: The Discovery of Quantum Mechanics*. New York: Springer-Verlag, 1982b.
- [Mehra and Rechenberg, 1982c] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 3: The formulation of matrix mechanics and its modifications, 1925-1926*. New York: Springer-Verlag, 1982c.
- [Mehra and Rechenberg, 1982d] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 4: The fundamental equations of quantum mechanics 1925-1926. The reception of the new quantum mechanics*. New York: Springer-Verlag, 1982d.
- [Mehra and Rechenberg, 1987] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 5: Erwin Schrödinger and the Rise of Wave Mechanics*. New York: Springer-Verlag, 1987.
- [Mehra and Rechenberg, 2000] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 6: The Completion of Quantum Mechanics 1926-1941. Part 1: The probabilistic Interpretation and the Empirical and Mathematical Foundation of Quantum Mechanics, 1926-1936*. New York: Springer-Verlag, 2000.
- [Mehra and Rechenberg, 2001] J. Mehra and H. Rechenberg. *The Historical Development of Quantum Theory. Vol. 6: The Completion of Quantum Mechanics 1926-1941. Part 2: The Conceptual Completion of Quantum Mechanics*. New York: Springer-Verlag, 2001.
- [Meinrenken, 1998] E. Meinrenken. Symplectic surgery and the Spin^c-Dirac operator. *Adv. Mathematical*, 134: 240-277, 1998.
- [Meinrenken and Sjamaar, 1999] E. Meinrenken and R. Sjamaar. Singular reduction and quantization. *Topology*, 38: 699-762, 1999.
- [Mermin, 2004] N. D. Mermin. What's wrong with this quantum world? *Physics Today*, 57 (2): 10, 2004.
- [Mielnik, 1968] B. Mielnik. Geometry of quantum states. *Communications in Mathematical Physics*, 9: 55-80, 1968.
- [Miller, 1984] A. I. Miller. *Imagery in Scientific Thought: Creating 20th-Century Physics*. Boston: Birkhäuser, 1984.
- [Mirlin, 2000] A. D. Mirlin. Statistics of energy levels and eigenfunctions in disordered systems. *Physics Reports*, 326: 259-382, 2000.
- [Mittelstaedt, 2004] P. Mittelstaedt. *The Interpretation of Quantum Mechanics and the Measurement Process*. Cambridge: Cambridge University Press, 2004.
- [Moore, 1939] G. E. Moore. Proof of an external world. *Proceedings of the British Academy*, 25: 273-300, 1939. Reprinted in *Philosophical Papers* (George, Allen and Unwin, London, 1959) and in *Selected Writings* (Routledge, London, 1993).
- [Moore, 1989] W. Moore. *Schrödinger: Life and Thought*. Cambridge: Cambridge University Press, 1989.
- [Morchio and Strocchi, 1987] G. Morchio and F. Strocchi. Mathematical structures for long-range dynamics and symmetry breaking. *Journal of Mathematical Physics*, 28: 622-635, 1987.
- [Muller, 1997] F. A. Muller. The equivalence myth of quantum mechanics I, II. *Studies in History and Philosophy of Modern Physics*, 28: 35-61, 219-247, 1997.
- [Murdoch, 1987] D. Murdoch. *Niels Bohrs Philosophy of Physics*. Cambridge: Cambridge University Press, 1987.
- [Nadirashvili et al., 2001] N. Nadirashvili, J. Toth, and D. Yakobson. Geometric properties of eigenfunctions. *Russian Mathematical Surveys*, 56: 1085-1105, 2001.
- [Nagy, 2000] G. Nagy. A deformation quantization procedure for C*-algebras. *Journal of Operator Theory*, 44: 369-411, 2000.
- [Narnhofer, 2001] H. Narnhofer. Quantum K-systems and their abelian models. *Foundations of Probability and Physics*, pages 274-302. River Edge, NJ: World Scientific, 2001.

- [Natsume and Nest, 1999] T. Natsume and R. Nest. Topological approach to quantum surfaces. *Communications in Mathematical Physics*, 202: 65–87, 1999.
- [Natsume *et al.*, 2003] T. Natsume, R. Nest, and P. Ingo. Strict quantizations of symplectic manifolds. *Letters in Mathematical Physics*, 66: 73–89, 2003.
- [Nauenberg, 1989] M. Nauenberg. Quantum wave packets on Kepler elliptic orbits. *Physical Review*, A40: 1133–1136, 1989.
- [Nauenberg *et al.*, 1994] M. Nauenberg, C. Stroud, and J. Yeazell. The classical limit of an atom. *Scientific American*, June, 24–29, 1994.
- [Neumann von, 1931] J. von Neumann. Die Eindeutigkeit der Schrödingerschen Operatoren. *Mathematische Annalen*, 104: 570–578, 1931.
- [Neumann von, 1932] J. von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Berlin: Springer-Verlag. English translation (1955): *Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press, 1932.
- [Neumann von, 1938] J. von Neumann. On infinite direct products. *Compositio Mathematica*, 6: 1–77, 1938.
- [Neumann von, 1981] J. von Neumann. Continuous geometries with a transition probability. *Memoirs of the American Mathematical Society*, 252: 1–210. (I.S. Halperin (ed.). MS from 1937), 1981.
- [Neumann, 1972] H. Neumann. Transformation properties of observables. *Helvetica Physica Acta*, 25: 811–819, 1972.
- [Nourrigat and Royer, 2004] J. Nourrigat and C. Royer. Thermodynamic limits for Hamiltonians defined as pseudodifferential operators. *Communications in Partial Differential Equations*, 29: 383–417, 2004.
- [O’Connor *et al.*, 1992] P. W. O’Connor, S. Tomsovic, and E. J. Heller. Semiclassical dynamics in the strongly chaotic regime: breaking the log time barrier. *Physica*, D55: 340–357, 1992.
- [Odzijewicz, 1992] A. Odziejewicz. Coherent states and geometric quantization. *Communications in Mathematical Physics* 150, 385–413, 1992.
- [Odzijewicz and Ratiu, 2003] A. Odziejewicz and T. S. Ratiu. Banach Lie-Poisson spaces and reduction. *Communications in Mathematical Physics*, 243: 1–54, 2003.
- [Olkiewicz, 1999a] R. Olkiewicz. Dynamical semigroups for interacting quantum and classical systems. *Journal of Mathematical Physics*, 40: 1300–1316, 1999a.
- [Olkiewicz, 1999b] R. Olkiewicz. Environment-induced superselection rules in Markovian regime. *Communications in Mathematical Physics*, 208: 245–265, 1999b.
- [Olkiewicz, 2000] R. Olkiewicz. Structure of the algebra of effective observables in quantum mechanics. *Annals of Physics*, 286: 10–22, 2000.
- [Ollivier *et al.*, 2004] H. Ollivier, D. Poulin, and W. H. Zurek. Environment as witness: selective proliferation of information and emergence of objectivity. 2004. [arXiv: quant-ph/0408125](https://arxiv.org/abs/quant-ph/0408125)
- [Olshanetsky and Perelomov, 1981] M. A. Olshanetsky and A. M. Perelomov. Classical integrable finite-dimensional systems related to Lie algebras. *Physics Reports*, 71: 313–400, 1981.
- [Olshanetsky and Perelomov, 1983] M. A. Olshanetsky and A. M. Perelomov. Quantum integrable systems related to Lie algebras. *Physics Reports*, 94: 313–404, 1983.
- [Omnès, 1992] R. Omnès. Consistent interpretations of quantum mechanics. *Reviews of Modern Physics*, 64: 339–382, 1992.
- [Omnès, 1994] R. Omnès. *The Interpretation of Quantum Mechanics*. Princeton: Princeton University Press, 1994.
- [Omnès, 1997] R. Omnès. Quantum-classical correspondence using projection operators. *Journal of Mathematical Physics*, 38: 697–707, 1997.
- [Omnès, 1999] R. Omnès. *Understanding Quantum Mechanics*. Princeton: Princeton University Press, 1999.
- [Ørsted, 1979] B. Ørsted. Induced representations and a new proof of the imprimitivity theorem. *Journal of Functional Analysis*, 31: 355–359, 1979.
- [Ozorio de Almeida, 1988] A. M. Ozorio de Almeida. *Hamiltonian Systems: Chaos and Quantization*. Cambridge: Cambridge University Press, 1988.
- [Pais, 1982] A. Pais. *Subtle is the Lord: The Science and Life of Albert Einstein*. Oxford: Oxford University Press, 1982.
- [Pais, 1991] A. Pais. *Niels Bohrs Times: In Physics, Philosophy, and Polity*. Oxford: Oxford University Press, 1991.
- [Pais, 1997] A. Pais. *A Tale of Two Continents: A Physicist’s Life in a Turbulent World*. Princeton: Princeton University Press, 1997.

- [Pais, 2000] A. Pais. *The Genius of Science*. Oxford: Oxford University Press, 2000.
- [Parthasarathy, 1992] K. R. Parthasarathy. *An Introduction to Quantum Stochastic Calculus*. Basel: Birkhäuser, 1992.
- [Paul and Uribe, 1995] T. Paul and A. Uribe. The semi-classical trace formula and propagation of wave packets. *Journal of Functional Analysis*, 132: 192–249, 1995.
- [Paul and Uribe, 1996] T. Paul and A. Uribe. On the pointwise behavior of semi-classical measures. *Communications in Mathematical Physics*, 175: 229–258, 1996.
- [Paul and Uribe, 1998] T. Paul and A. Uribe. A. Weighted trace formula near a hyperbolic trajectory and complex orbits. *Journal of Mathematical Physics*, 39: 4009–4015, 1998.
- [Pauli, 1925] W. Pauli. Über den Einfluß der Geschwindigkeitsabhängigkeit der Elektronenmasse auf den Zeemaneffekt. *Zeitschrift für Physik*, 31: 373–385, 1925.
- [Pauli, 1933] W. Pauli. *Die allgemeinen Prinzipien der Wellenmechanik*. Flügge, S. (Ed.). *Handbuch der Physik*, Vol. V, Part I, 1933. Translated as W. Pauli, *General Principles of Quantum Mechanics*. Berlin: Springer-Verlag, 1980.
- [Pauli, 1949] W. Pauli. Die philosophische Bedeutung der Idee der Komplementarität. 1949. Reprinted in K. von Meyenn (ed.), *Wolfgang Pauli: Physik und Erkenntnistheorie*, pages 10–23. Braunschweig: Vieweg Verlag, 1984. English translation in Pauli, 1994.
- [Pauli, 1979] W. Pauli. *Wissenschaftlicher Briefwechsel mit Bohr, Einstein, Heisenberg. Vol 1: 1919–1929*. A. Hermann, K. von Meyenn, and V. Weisskopf (eds.), New York: Springer-Verlag, 1979.
- [Pauli, 1985] W. Pauli. *Wissenschaftlicher Briefwechsel mit Bohr, Einstein, Heisenberg. Vol 2: 1930–1939*. K. von Meyenn (ed.). New York: Springer-Verlag, 1985.
- [Pauli, 1994] W. Pauli. *Writings on Physics and Philosophy*. C. P. Enz and K. von Meyenn (eds.). Berlin: Springer-Verlag, 1994.
- [Paz and Zurek, 1999] J. P. Paz and W. H. Zurek. Quantum limit of decoherence: environment induced superselection of energy eigenstates. *Physical Review Letters*, 82: 5181–5185, 1999.
- [Pedersen, 1979] G. K. Pedersen. *C*-algebras and their Automorphism Groups*. London: Academic Press, 1979.
- [Pedersen, 1989] G. K. Pedersen. *Analysis Now*. New York: Springer-Verlag, 1989.
- [Perelomov, 1986] A. Perelomov. *Generalized Coherent States and their Applications*. Berlin: Springer-Verlag, 1986.
- [Peres, 1984] A. Peres. Stability of quantum motion in chaotic and regular systems. *Physical Review*, A30: 1610–1615, 1984.
- [Peres, 1995] A. Peres. *Quantum Theory: Concepts and Methods*. Dordrecht: Kluwer Academic Publishers, 1995.
- [Petersen, 1963] A. Petersen. The Philosophy of Niels Bohr. *Bulletin of the Atomic Scientists*, 19: 8–14, 1963.
- [Pitowsky, 1989] I. Pitowsky. *Quantum Probability - Quantum Logic*. Berlin: Springer-Verlag, 1989.
- [Planck, 1906] M. Planck. *Vorlesungen Über die Theorie der Wärmestrahlung* Leipzig: J.A. Barth, 1906.
- [Poincaré, 1892–1899] H. Poincaré. *Les Méthodes Nouvelles de la Mécanique Céleste*. Paris: Gauthier-Villars, 1892–1899.
- [Popov, 2000] G. Popov. Invariant tori, effective stability, and quasimodes with exponentially small error terms. I and II. *Annales Henri Poincaré*, 1: 223–248 and 249–279, 2000.
- [Poulin, 2004] D. Poulin. Macroscopic observables. 2004. [arXiv:quant-ph/0403212](https://arxiv.org/abs/quant-ph/0403212).
- [Poulsen, 1970] N. S. Poulsen. *Regularity Aspects of the Theory of Infinite-Dimensional Representations of Lie Groups*. Ph.D Thesis, MIT, 1970.
- [Primas, 1983] H. Primas. *Chemistry, Quantum Mechanics and Reductionism*. Second Edition. Berlin: Springer-Verlag, 1983.
- [Primas, 1997] H. Primas. The representation of facts in physical theories. In H. Atmanspacher and E. Ruhnau (eds.), *Time, Temporality, Now*, pages 241–263. Berlin: Springer-Verlag, 1997.
- [Prugovecki, 1971] E. Prugovecki. *Quantum Mechanics in Hilbert Space*. New York: Academic Press, 1971.
- [Putza, 1993] M. Putza. *Hamiltonian Dynamical Systems and Geometric Quantization*. Dordrecht: D. Reidel, 1993.
- [Raggio, 1981] G. A. Raggio. *States and Composite Systems in W*-algebras Quantum Mechanics*. Ph.D Thesis, ETH Zürich, 1981.

- [Raggio, 1988] G. A. Raggio. A remark on Bell's inequality and decomposable normal states. *Letters in Mathematical Physics*, 15: 27–29, 1988.
- [Raggio and Werner, 1989] G. A. Raggio and R. F. Werner. Quantum statistical mechanics of general mean field systems. *Helvetica Physica Acta*, 62: 980–1003, 1989.
- [Raggio and Werner, 1991] G. A. Raggio and R. F. Werner. The Gibbs variational principle for inhomogeneous mean field systems. *Helvetica Physica Acta*, 64: 633–667, 1991.
- [Raimond *et al.*, 2001] R. M. Raimond, M. Brune, and S. Haroche. Manipulating quantum entanglement with atoms and photons in a cavity. *Reviews of Modern Physics*, 73: 565–582, 2001.
- [Rédei, 1998] M. Rédei. *Quantum logic in algebraic approach*. Dordrecht: Kluwer Academic Publishers, 1998.
- [Rédei and Stöltzner, 2001] M. Rédei and M. Stöltzner (eds.). *John von Neumann and the Foundations of Modern Physics*. Dordrecht: Kluwer Academic Publishers, 2001.
- [Reed and Simon, 1972] M. Reed and B. Simon. *Methods of Modern Mathematical Physics. Vol I. Functional Analysis*. New York: Academic Press, 1972.
- [Reed and Simon, 1975] M. Reed and B. Simon. *Methods of Modern Mathematical Physics. Vol II. Fourier Analysis, Self-adjointness*. New York: Academic Press, 1975.
- [Reed and Simon, 1978] M. Reed and B. Simon. *Methods of Modern Mathematical Physics. Vol IV. Analysis of Operators*. New York: Academic Press, 1978.
- [Reed and Simon, 1979] M. Reed and B. Simon. *Methods of Modern Mathematical Physics. Vol III. Scattering Theory*. New York: Academic Press, 1979.
- [Reichl, 2004] L. E. Reichl. *The Transition to Chaos in Conservative Classical Systems: Quantum Manifestations*. Second Edition. New York: Springer-Verlag, 2004.
- [Rieckers, 1984] A. Rieckers. On the classical part of the mean field dynamics for quantum lattice systems in grand canonical representations. *Journal of Mathematical Physics*, 25: 2593–2601, 1984.
- [Rieffel, 1989a] M. A. Rieffel. Deformation quantization of Heisenberg manifolds. *Communications in Mathematical Physics* 122, 531–562, 1989a.
- [Rieffel, 1989b] M. A. Rieffel. Continuous fields of C^* -algebras coming from group cocycles and actions. *Mathematical Annals*, 283: 631–643, 1989b.
- [Rieffel, 1994] M. A. Rieffel. Quantization and C^* -algebras. *Contemporary Mathematics*, 167: 66–97, 1994.
- [Riesz and Sz.-Nagy, 1990] F. Riesz and B. Sz.-Nagy. *Functional Analysis*. New York: Dover, 1990.
- [Robert, 1987] D. Robert. *Autour de l'Approximation Semi-Classique*. Basel: Birkhäuser, 1987.
- [Robert, 1992] D. Robert (ed.). *Méthodes Semi-Classiques. Astérisque*, 207: 1–212, *ibid.*: 210, 1–384, 1992.
- [Robert, 1998] D. Robert. Semi-classical approximation in quantum mechanics. A survey of old and recent mathematical results. *Helvetica Physica Acta*, 71: 44–116, 1998.
- [Roberts, 1990] J. E. Roberts. Lectures on algebraic quantum field theory. In D. Kastler (ed.), *The Algebraic Theory of Superselection Sectors. Introduction and Recent Results*, pages 1–112. River Edge, NJ: World Scientific Publishing Co, 1990.
- [Roberts and Roepstorff, 1969] J. E. Roberts and G. Roepstorff. Some basic concepts of algebraic quantum theory. *Communications in Mathematical Physics*, 11: 321–338, 1969.
- [Robinett, 2004] R. W. Robinett. Quantum wave packet revival. *Physics Reports*, 392: 1–119, 2004.
- [Robinson, 1994] D. Robinson. Can superselection rules solve the measurement problem? *British Journal for the Philosophy of Science*, 45: 79–93, 1994.
- [Robinson, 1988a] S. L. Robinson. The semiclassical limit of quantum mechanics. I. Time evolution. *Journal of Mathematical Physics*, 29: 412–419, 1988a.
- [Robinson, 1988b] S. L. Robinson. The semiclassical limit of quantum mechanics. II. Scattering theory. *Annales de l'Institut Henri Poincaré*, A48: 281–296, 1988b.
- [Robinson, 1993] S. L. Robinson. Semiclassical mechanics for time-dependent Wigner functions. *Journal of Mathematical Physics*, 34: 2185–2205, 1993.
- [Robson, 1996] M. A. Robson. Geometric quantization of reduced cotangent bundles. *Journal of Geometry and Physics*, 19: 207–245, 1996.
- [Rosenfeld, 1967] L. Rosenfeld. Niels Bohr in the Thirties. Consolidation and extension of the conception of complementarity. In S. Rozental (ed.), *Niels Bohr: His Life and Work as Seen by His Friends and Colleagues*, pages 114–136. Amsterdam: North-Holland, 1967.

- [Rudolph, 1996a] O. Rudolph. Consistent histories and operational quantum theory. *International Journal of Theoretical Physics*, 35: 1581–1636, 1996a.
- [Rudolph, 1996b] O. Rudolph. On the consistent effect histories approach to quantum mechanics. *Journal of Mathematical Physics*, 37: 5368–5379, 1996b.
- [Rudolph, 2000] O. Rudolph. The representation theory of decoherence functionals in history quantum theories. *International Journal of Theoretical Physics*, 39: 871–884, 2000.
- [Rudolph and Wright, 1999] O. Rudolph and J. D. M. Wright. Homogeneous decoherence functionals in standard and history quantum mechanics. *Communications in Mathematical Physics*, 204: 249–267, 1999.
- [Sarnak, 1999] P. Sarnak. Quantum chaos, symmetry and zeta functions. I. and II. *Quantum Chaos. Current Developments in Mathematics, 1997 (Cambridge, MA)*, pages 127–144 and 145–159. Boston: International Press, 1999.
- [Saunders, 1993] S. Saunders. Decoherence, relative states, and evolutionary adaptation. *Foundations of Physics*, 23: 1553–1585, 1993.
- [Saunders, 1995] S. Saunders. Time, quantum mechanics, and decoherence. *Synthese*, 102: 235–266, 1995.
- [Saunders, 2005] S. Saunders. Complementarity and scientific rationality. *Foundations of Physics*, 35:417–447, 2005.
- [Scheibe, 1973] E. Scheibe. *The Logical Analysis of Quantum Mechanics*. Oxford: Pergamon Press, 1973.
- [Scheibe, 1991] E. Scheibe. J. v. Neumanns und J. S. Bells Theorem. Ein Vergleich. *Philosophia Naturalis*, 28: 35–53, 1991. English translation in E. Scheibe, *Between Rationalism and Empiricism: Selected Papers in the Philosophy of Physics*. New York: Springer-Verlag, 2001.
- [Scheibe, 1999] E. Scheibe. *Die Reduktion Physikalischer Theorien. Ein Beitrag zur Einheit der Physik. Teil II: Inkommensurabilität und Grenzfallreduktion*. Berlin: Springer-Verlag, 1999.
- [Schlosshauer, 2004] M. Schlosshauer. Decoherence, the measurement problem, and interpretations of quantum mechanics. *Reviews of Modern Physics*, 76: 1267–1306, 2004.
- [Schmüdgen, 1990] K. Schmüdgen. *Unbounded Operator Algebras and Representation Theory*. Basel: Birkhäuser Verlag, 1990.
- [Schrödinger, 1926a] E. Schrödinger. Quantisierung als Eigenwertproblem. I–IV. *Annalen der Physik*, 79: 361–76, 489–527, *ibid.* 80: 437–90, *ibid.* 81: 109–39, 1926a. English translation in Schrödinger (1928).
- [Schrödinger, 1926b] E. Schrödinger. Der stetige Übergang von der Mikro-zur Makromechanik. *Die Naturwissenschaften*, 14: 664–668, 1926b. English translation in Schrödinger (1928).
- [Schrödinger, 1926c] E. Schrödinger. Über das Verhältnis der Heisenberg–Born–Jordanschen Quantenmechanik zu der meinen. *Annalen der Physik*, 79: 734–56, 1926c. English translation in Schrödinger (1928).
- [Schrödinger, 1928] E. Schrödinger. *Collected Papers on Wave Mechanics*. London: Blackie and Son, 1928.
- [Schroek, 1996] F. E. Schroek Jr.. *Quantum Mechanics on Phase Space*. Dordrecht: Kluwer Academic Publishers, 1996.
- [Scutaru, 1977] H. Scutaru. Coherent states and induced representations. *Letters in Mathematical Physics*, 2: 101–107, 1977.
- [Segal, 1960] I. E. Segal. Quantization of nonlinear systems. *Journal of Mathematical Physics*, 1: 468–488, 1960.
- [Sewell, 1986] G. L. Sewell. *Quantum Theory of Collective Phenomena*. New York: Oxford University Press, 1986.
- [Sewell, 2002] G. L. Sewell. *Quantum Mechanics and its Emergent Macrophysics*. Princeton: Princeton University Press, 2002.
- [Simon, 1976] B. Simon. Quantum dynamics: from automorphism to Hamiltonian. *Studies in Mathematical Physics: Essays in Honour of Valentine Bargmann*, pp. 327–350. Lieb, E.H., Simon, B. and Wightman, A.S. (Eds.). Princeton: Princeton University Press, 1976.
- [Simon, 1980] B. Simon. The classical limit of quantum partition functions. *Communications in Mathematical Physics*, 71: 247–276, 1980.
- [Simon, 2000] B. Simon. Schrödinger operators in the twentieth century. *Journal of Mathematical Physics*, 41: 3523–3555, 2000.
- [Śniatycki, 1980] J. Śniatycki. *Geometric Quantization and Quantum Mechanics*. Berlin: Springer-Verlag, 1980.

- [Snirelman, 1974] A. I. Snirelman. Ergodic properties of eigenfunctions. *Uspekhi Mathematical Nauk*, 29: 181–182, 1974.
- [Souriau, 1969] J.-M. Souriau. *Structure des systèmes dynamiques*. Paris: Dunod, 1969. Translated as J.-M. Souriau (1997).
- [Souriau, 1997] J.-M. Souriau. *Structure of Dynamical Systems. A Symplectic View of Physics*. Boston: Birkhäuser, 1997.
- [Stapp, 1972] H. P. Stapp. The Copenhagen Interpretation. *American Journal of Physics*, 40: 1098–1116, 1972.
- [Steiner, 1998] M. Steiner. *The Applicability of Mathematics as a Philosophical Problem*. Cambridge (MA): Harvard University Press, 1998.
- [Stinespring, 1955] W. Stinespring. Positive functions on C^* -algebras. *Proceedings of the American Mathematical Society*, 6: 211–216, 1955.
- [Størmer, 1969] E. Størmer. Symmetric states of infinite tensor products of C^* -algebras. *Journal of Functional Analysis*, 3: 48–68, 1969.
- [Strawson, 1959] P. F. Strawson. *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen, 1959.
- [Streater, 2000] R. F. Streater. Classical and quantum probability. *Journal of Mathematical Physics*, 41: 3556–3603, 2000.
- [Strichartz, 1983] R. S. Strichartz. Analysis of the Laplacian on a complete Riemannian manifold. *Journal of Functional Analysis*, 52: 48–79, 1983.
- [Strocchi, 1985] F. Strocchi. *Elements of Quantum mechanics of Infinite Systems*. Singapore: World Scientific, 1985.
- [Strunz et al., 2003] W. T. Strunz, F. Haake, and D. Braun. Universality of decoherence for macroscopic quantum superpositions. *Physical Review*, A 67: 022101–022114, 2003.
- [Summers and Werner, 1987] S. J. Summers and R. Werner. Bells inequalities and quantum field theory, I, II. *Journal of Mathematical Physics*, 28: 2440–2447, 2448–2456, 1987.
- [Sundermeyer, 1982] K. Sundermeyer. *Constrained Dynamics*. Berlin: Springer-Verlag, 1982.
- [Takesaki, 2003] M. Takesaki. *Theory of Operator Algebras. Vols. I-III*. New York: Springer-Verlag, 2003.
- [Thirring and Wehrl, 1967] W. Thirring and A. Wehrl. On the mathematical structure of the BCS model. I *Communications in Mathematical Physics*, 4: 303–314, 1967.
- [Thirring, 1968] W. Thirring. On the mathematical structure of the BCS model. II. *Communications in Mathematical Physics*, 7: 181–199, 1968.
- [Thirring, 1981] W. Thirring. *A Course in Mathematical Physics. Vol. 3: Quantum Mechanics of Atoms and Molecules*. New York: Springer-Verlag, 1981.
- [Thirring, 1983] W. Thirring. *A Course in Mathematical Physics. Vol. 4: Quantum Mechanics of Large Systems*. New York: Springer-Verlag, 1983.
- [Tomsovic and Heller, 1993] S. Tomsovic and E. J. Heller. Long-time semiclassical dynamics of chaos: The stadium billiard. *Physical Review*, E47: 282–299.
- [Tomsovic and Heller, 2002] S. Tomsovic and E. J. Heller. Comment on “Ehrenfest times for classically chaotic systems”. *Physical Review*, E65, 035208-1–2, 2002.
- [Toth, 1996] J. A. Toth. Eigenfunction localization in the quantized rigid body. *Journal of Differential Geometry*, 43: 844–858, 1996.
- [Toth, 1999] J. A. Toth. On the small-scale mass concentration of modes. *Communications in Mathematical Physics*, 206: 409–428, 1999.
- [Toth and Zelditch, 2002] J. A. Toth and S. Zelditch. Riemannian manifolds with uniformly bounded eigenfunctions. *Duke Mathematical Journal*, 111: 97–132, 2002.
- [Toth and Zelditch, 2003a] J. A. Toth and S. Zelditch. L^p norms of eigenfunctions in the completely integrable case. *Annales Henri Poincaré*, 4: 343–368, 2003a.
- [Toth and Zelditch, 2003b] J. A. Toth and S. Zelditch. Norms of modes and quasi-modes revisited. *Contemporary Mathematics*, 320: 435–458, 2003b.
- [Tuynman, 1987] G. M. Tuynman. Quantization: towards a comparison between methods. *Journal of Mathematical Physics*, 28: 2829–2840, 1987.
- [Tuynman, 1998] G. M. Tuynman. Prequantization is irreducible. *Indagationes Mathematicae (New Series)*, 9: 607–618, 1998.
- [Unnerstall, 1990a] T. Unnerstall. Phase-spaces and dynamical descriptions of infinite mean-field quantum systems. *Journal of Mathematical Physics*, 31: 680–688, 1990a.
- [Unnerstall, 1990b] T. Unnerstall. Schrödinger dynamics and physical folia of infinite mean-field quantum systems. *Communications in Mathematical Physics*, 130: 237–255, 1990b.

- [Vaisman, 1991] I. Vaisman, I. On the geometric quantization of Poisson manifolds. *Journal of Mathematical Physics*, 32: 3339–3345, 1991.
- [van Fraassen, 1991] B. C. van Fraassen. *Quantum Mechanics: An Empiricist View*. Oxford: Oxford University Press, 1991.
- [van Hove, 1951] L. van Hove. Sur certaines représentations unitaires d'un groupe infini de transformations. *Memoires de l'Académie Royale de Belgique, Classe des Sciences*, 26: 61–102, 1951.
- [Van Vleck, 1928] J. H. Van Vleck. The Correspondence Principle in the Statistical Interpretation of Quantum Mechanics. *Proceedings of the National Academy of Sciences*, 14: 178–188, 1928.
- [van der Waerden, 1967] B. L. van der Waerden (ed.). *Sources of Quantum Mechanics*. Amsterdam: North-Holland, 1967.
- [van Kampen, 1954] N. van Kampen. Quantum statistics of irreversible processes. *Physica* 20, 603–622, 1954.
- [van Kampen, 1988] N. van Kampen. Ten theorems about quantum mechanical measurements. *Physica*, A153: 97–113, 1988.
- [van Kampen, 1993] N. van Kampen. Macroscopic systems in quantum mechanics. *Physica*, A194: 542–550, 1993.
- [Vanicek and Heller, 2003] J. Vanicek and E. J. Heller. Semiclassical evaluation of quantum fidelity *Physical Review*, E68, 056208-1–5, 2003.
- [Vergne, 1994] M. Vergne. Geometric quantization and equivariant cohomology. *First European Congress in Mathematics, Vol. 1*, pages 249–295. Boston: Birkhäuser, 1994.
- [Vermaas, 2000] P. Vermaas. *A Philosopher's Understanding of Quantum Mechanics: Possibilities and Impossibilities of a Modal Interpretation*. Cambridge: Cambridge University Press, 2000.
- [Vey, 1975] J. Vey. Déformation du crochet de Poisson sur une variété symplectique. *Commentarii Mathematici Helvetici*, 50: 421–454, 1975.
- [Voros, 1979] A. Voros. Semi-classical ergodicity of quantum eigenstates in the Wigner representation. *Stochastic Behaviour in Classical and Quantum Hamiltonian Systems. Lecture Notes in Physics*, 93: 326–333, 1979.
- [Wallace, 2002] D. Wallace. Worlds in the Everett interpretation. *Studies in History and Philosophy of Modern Physics*, 33B: 637–661, 2002.
- [Wallace, 2003] D. Wallace. Everett and structure. *Studies in History and Philosophy of Modern Physics*, 34B: 87–105, 2003.
- [Wan and Fountain, 1998] K. K. Wan and R. H. Fountain. Quantization by parts, maximal symmetric operators, and quantum circuits. *International Journal of Theoretical Physics*, 37: 2153–2186, 1998.
- [Wan et al., 1999] K. K. Wan, J. Bradshaw, C. Trueman, and F. E. Harrison. Classical systems, standard quantum systems, and mixed quantum systems in Hilbert space. *Foundations of Physics*, 28: 1739–1783, 1998.
- [Wang, 1986] X.-P. Wang. Approximation semi-classique de l'équation de Heisenberg. *Communications in Mathematical Physics*, 104: 77–86, 1986.
- [Wegge-Olsen, 1993] N. E. Wegge-Olsen. *K-theory and C*-algebras*. Oxford: Oxford University Press, 1993.
- [Weinstein, 1983] A. Weinstein. The local structure of Poisson manifolds. *Journal of Differential Geometry*, 18: 523–557, 1983.
- [Werner, 1995] R. F. Werner. The classical limit of quantum theory. 1995. [arXiv:quant-ph/9504016](https://arxiv.org/abs/quant-ph/9504016).
- [Werner, 1989] R. F. Werner. Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model. *Physical Review*, A40: 42774281, 1989.
- [Weyl, 1931] H. Weyl. *The Theory of Groups and Quantum Mechanics*. New York: Dover, 1931.
- [Wheeler and Zurek, 1983] J. A. Wheeler and W. H. Zurek (eds.). *Quantum Theory and Measurement*. Princeton: Princeton University Press, 1983.
- [Whitten-Wolfe and Emch, 1976] B. Whitten-Wolfe and G. G. Emch. A mechanical quantum measuring process. *Helvetica Physica Acta*, 49: 45–55, 1976.
- [Wick et al., 1952] C. G. Wick, A. S. Wightman, and E. P. Wigner. The intrinsic parity of elementary particles. *Physical Review*, 88: 101–105, 1952.
- [Wightman, 1962] A. S. Wightman. On the localizability of quantum mechanical systems. *Reviews of Modern Physics*, 34: 845–872, 1962.

- [Wigner, 1932] E. P. Wigner. On the quantum correction for thermodynamic equilibrium. *Physical Review*, 40: 749–759, 1932.
- [Wigner, 1939] E. P. Wigner. Unitary representations of the inhomogeneous Lorentz group. *Annals of Mathematics*, 40: 149–204, 1939.
- [Wigner, 1963] E. P. Wigner. The problem of measurement. *American Journal of Physics*, 31: 6–15, 1963.
- [Williams, 2003] F. L. Williams. *Topics in Quantum Mechanics*. Boston: Birkhäuser, 2003.
- [Woodhouse, 1992] N. M. J. Woodhouse. *Geometric Quantization*. Second edition. Oxford: The Clarendon Press, 1992.
- [Yajima, 1979] K. Yajima. The quasi-classical limit of quantum scattering theory. *Communications in Mathematical Physics*, 69: 101–129, 1979.
- [Zaslavsky, 1981] G. M. Zaslavsky. Stochasticity in quantum systems. *Physics Reports*, 80: 157–250, 1981.
- [Zeh, 1970] H. D. Zeh. On the interpretation of measurement in quantum theory. *Foundations of Physics*, 1: 69–76, 1970.
- [Zelditch, 1987] S. Zelditch. Uniform distribution of eigenfunctions on compact hyperbolic surfaces. *Duke Mathematical J.*, 55: 919–941, 1987.
- [Zelditch, 1990] S. Zelditch. Quantum transition amplitudes for ergodic and for completely integrable systems. *Journal of Functional Analysis*, 94: 415–436, 1990.
- [Zelditch, 1991] S. Zelditch. Mean Lindelöf hypothesis and equidistribution of cusp forms and Eisenstein series. *Journal of Functional Analysis*, 97, 1–49, 1991.
- [Zelditch, 1992a] S. Zelditch. On a “quantum chaos” theorem of R. Schrader and M. Taylor. *Journal of Functional Analysis*, 109: 1–21, 1992a.
- [Zelditch, 1992b] S. Zelditch. Quantum ergodicity on the sphere. *Communications in Mathematical Physics*, 146: 61–71, 1992b.
- [Zelditch, 1996a] S. Zelditch. Quantum dynamics from the semiclassical viewpoint. Lectures at the Centre E. Borel, 1996a. Available at <http://mathnt.mat.jhu.edu/zelditch/Preprints/preprints.html>.
- [Zelditch, 1996b] S. Zelditch. Quantum mixing. *Journal of Functional Analysis*, 140: 68–86, 1996b.
- [Zelditch, 1996c] S. Zelditch. Quantum ergodicity of C^* dynamical systems. *Communications in Mathematical Physics*, 177: 507–528, 1996c.
- [Zelditch and Zworski, 1996] S. Zelditch and M. Zworski. Ergodicity of eigenfunctions for ergodic billiards. *Communications in Mathematical Physics*, 175: 673–682, 1996.
- [Zurek, 1981] W. H. Zurek. Pointer basis of quantum apparatus: into what mixture does the wave packet collapse? *Physical Review*, D24: 1516–1525, 1981.
- [Zurek, 1982] W. H. Zurek. Environment-induced superselections rules. *Physical Review*, D26: 1862–1880, 1982.
- [Zurek, 1991] W. H. Zurek. Decoherence and the transition from quantum to classical. *Physics Today*, 44 (10): 36–44, 1991.
- [Zurek, 1993] W. H. Zurek. Negotiating the tricky border between quantum and classical. *Physics Today*, 46 (4): 13–15, 81–90, 1993.
- [Zurek, 2003] W. H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75: 715–775, 2003.
- [Zurek, 2005] W. H. Zurek. Probabilities from entanglement, Born’s rule $p_k = |\psi|^2$ from environment. *Physical Review*, A 71, 052105, 2005.
- [Zurek et al., v] W. H. Zurek, S. Habib, and J. P. Paz. Coherent states via decoherence. *Physical Review Letters*, 70: 1187–1190, 1993.
- [Zurek and Paz, 1995] W. H. Zurek and J. P. Paz. Why we don’t need quantum planetary dynamics: decoherence and the correspondence principle for chaotic systems. In D. H. Feng et al. (eds.), *Proceedings of the Fourth Drexel Meeting*. New York: Plenum, 1995. [arXiv:quant-ph/9612037](https://arxiv.org/abs/quant-ph/9612037)

This page intentionally left blank

QUANTUM INFORMATION AND COMPUTATION

Jeffrey Bub

1 INTRODUCTION

The subject of quantum information has its roots in the debate about conceptual issues in the foundations of quantum mechanics.

The story really begins with the dispute between Einstein and Bohr about the interpretation of quantum states, in particular the interpretation of so-called ‘entangled states’, which exhibit peculiar nonlocal statistical correlations for widely separated quantum systems. See, for example, [Bohr, 1949, 283] and Einstein’s reply in the same volume [Schilpp, 1949]. Einstein took the position that quantum mechanics is simply an incomplete theory. On the basis of a certain restricted set of correlations for a pair of systems in a particular entangled state, Einstein, Podolsky, and Rosen (EPR) argued in a seminal paper [Einstein *et al.*, 1935] that the phenomenon of entanglement conflicts with certain basic realist principles of separability and locality that all physical theories should respect, unless we regard quantum states as incomplete descriptions.

Bohr’s view, which he termed ‘complementarity’, eventually became entrenched as the orthodox Copenhagen interpretation, a patchwork of reformulations by Heisenberg, Pauli, von Neumann, Dirac, Wheeler, and others. (For a discussion, see [Howard, 2004] and Landsman, this vol., ch. 5.) As Pauli put it in correspondence with Max Born [Born, 1971, 218], a ‘detached observer’ description of the sort provided by classical physics is precluded by the nature of quantum phenomena, and a quantum description of events is as complete as it can be (in principle). Any application of quantum theory requires a ‘cut’ between the observer and the observed, or the macroscopic measuring instrument and the measured system, so that the description is in a certain sense contextual, where the relevant context is defined by the whole macroscopic experimental arrangement. So, for example, a ‘position measurement context’ provides information about position but excludes, in principle, the possibility of simultaneously obtaining momentum information, because there is no fact of the matter about momentum in this context: the momentum value is indeterminate. The Copenhagen interpretation conflicts with Einstein’s realism, his ‘philosophical prejudice’, as Pauli characterized it in a letter to Born [Born, 1971, 221], that lies at the heart of the dispute between Einstein and Bohr about the significance of the transition from classical to quantum mechanics.

The 1990's saw the development of a quantum theory of information, based on the realization that entanglement, rather than being a minor source of embarrassment for physics that need only concern philosophers, can actually be exploited as a nonclassical communication channel to perform information-processing tasks that would be impossible in a classical world. In a two-part commentary on the EPR paper, Schrödinger [1935, 555] identified entanglement as '*the* characteristic trait of quantum theory, the one that enforces its entire departure from classical lines of thought.' This has led to an explosive surge of research among physicists and computer scientists on the application of information-theoretic ideas to quantum computation (which exploits entanglement in the design of a quantum computer, so as to enable the efficient performance of certain computational tasks), to quantum communication (new forms of 'entanglement-assisted' communication, such as quantum teleportation), and to quantum cryptography (the identification of cryptographic protocols that are guaranteed to be unconditionally secure against eavesdropping or cheating, by the laws of quantum mechanics, even if all parties have access to quantum computers).

Some milestones: Bell's analysis [1964] turned the EPR argument on its head by showing that Einstein's assumptions of separability and locality, applicable in classical physics and underlying the EPR incompleteness argument, are incompatible with certain quantum statistical correlations (not explicitly considered by EPR) of separated systems in EPR-type entangled states. Later experiments [Aspect *et al.*, 1981; Aspect *et al.*, 1982] confirmed these nonclassical correlations in set-ups that excluded the possibility of any sort of physically plausible, non-superluminal, classical communication between the separated systems.

In the 1980s, various authors, e.g., Wiesner, Bennett, and Brassard [Wiesner, 1983; Bennett and Brassard, 1984; Bennett *et al.*, 1982] pointed out that one could exploit features of the measurement process in quantum mechanics to thwart the possibility of undetected eavesdropping in certain cryptographic procedures, specifically in key distribution — a procedure where two parties, Alice and Bob, who initially share no information end up each holding a secret random key which can be used to send encrypted messages between them. No third party, Eve, can obtain any information about the communications between Alice and Bob that led to the establishment of the key, without Alice and Bob becoming aware of Eve's interference, because Eve's measurements necessarily disturb the quantum states of the systems in the communication channel.

Bennett [1973] showed how to make a universal Turing machine reversible for any computation, a required step in the design of a quantum computer that evolves via unitary (and hence reversible) state transformations, and Benioff [1980] developed Hamiltonian models for computer computers. Feynman [1982] considered the problem of efficiently simulating the evolution of physical systems using quantum resources (noting that the classical simulation of a quantum process would be exponentially costly), which involves the idea of a quantum computation, but it was Deutsch [1985; 1989] who characterized the essential features of a universal quantum computer and formulated the first genuinely quantum algorithm.

Following Duetsch's work on quantum logic gates and quantum networks, several quantum algorithms were proposed for performing computational tasks more efficiently than any known classical algorithm, or in some cases more efficiently than any classical algorithm. The most spectacular of these is Shor's algorithm [1994; 1997] for finding the two prime factors of a positive integer $N = pq$, which is exponentially faster than the best-known classical algorithm. Since prime factorization is the basis of the most widely used public key encryption scheme (currently universally applied in communications between banks and commercial transactions over the internet), Shor's result has enormous practical significance.

In the following, I present an account of some of the theoretical developments in quantum information, quantum communication, quantum cryptography, and quantum computation. I conclude by considering whether a perspective in terms of quantum information suggests a new way of resolving the foundational problems of quantum mechanics that were the focus of the debate between Einstein and Bohr.

My discussion is heavily indebted to Michael Nielsen and Isaac Chuang's illuminating and comprehensive *Quantum Computation and Quantum Information* [2000], and to several insightful review articles: 'The Joy of Entanglement' by Sandu Popescu and Daniel Rohrlich [1998], 'Quantum Information and its Properties' by Richard Jozsa [1998], and 'Quantum Computing' by Andrew Steane [1998].

2 CLASSICAL INFORMATION

2.1 *Classical Information Compression and Shannon Entropy*

In this section, I review the basic elements of classical information theory. In §2.1, I introduce the notion of the Shannon entropy of an information source and the fundamental idea of information compression in Shannon's source coding theorem (or noiseless channel coding theorem). In §2.2, I define some information-theoretic concepts relevant to Shannon's noisy channel coding theorem.

The classical theory of information was initially developed to deal with certain questions in the communication of electrical signals. Shannon's ground-breaking paper 'A Mathematical Theory of Communication' [Shannon, 1948] followed earlier work by people like Nyquist [1924] and Hartley [1928] in the 1920s. The basic problem was the representation of messages, selected from an ensemble generated by a stochastic process at the message source, in such a way as to ensure their efficient transmission over an electrical circuit such as a noisy telegraph wire.

A communication set-up involves a transmitter or source of information, a (possibly noisy) channel, and a receiver. The source produces messages in the form of sequences of symbols from some alphabet, which Shannon represented mathematically as sequences of values of independent, identically distributed random variables. In later idealizations, the source is represented as *stationary*, in the sense (roughly) that the probability of any symbol (or n -tuple of symbols) appearing at

any given position in a (very long) sequence, when that position is considered with respect to an ensemble of possible sequences, is the same for all positions in the sequence, and *ergodic*, in the sense that this ‘ensemble average’ probability is equal to the ‘time average’ probability, where the time average refers to the probability of a symbol (or n -tuple of symbols) in a given (very long) sequence.

The fundamental question considered by Shannon was how to quantify the minimal physical resources required to store messages produced by a source, so that they could be communicated via a channel without loss and reconstructed by a receiver. Shannon’s source coding theorem (or noiseless channel coding theorem) answers this question.

To see the idea behind the theorem, consider a source that produces long sequences (messages) composed of symbols from a finite alphabet a_1, a_2, \dots, a_k , where the individual symbols are produced with probabilities p_1, p_2, \dots, p_k . A given sequence of symbols is represented as a sequence of values of independent, identically distributed, discrete random variables X_1, X_2, \dots . A *typical* sequence of length n , for large n , will contain close to $p_i n$ symbols a_i , for $i = 1, \dots, k$. So the probability of a sufficiently long typical sequence (assuming independence) will be:

$$(1) \quad p(x_1, x_2, \dots, x_n) = p(x_1)p(x_2) \dots p(x_n) \approx p_1^{p_1 n} p_2^{p_2 n} \dots p_k^{p_k n}.$$

Taking the logarithm of both sides (conventionally, in information theory, to the base 2) yields:

$$(2) \quad \log p(x_1, \dots, x_n) \approx n \sum_i p_i \log p_i := -nH(X)$$

where $H(X) := -\sum_i p_i \log p_i$ is the *Shannon entropy* of the source.

We can think about information in Shannon’s sense in various ways. We can take $-\log p_i$, a decreasing function of p_i with a minimum value of 0 when $p_i = 1$ for some i , as a measure of the information associated with identifying the symbol a_i produced by an information source. Then $H(X) = -\sum_i p_i \log p_i$ is the average information gain, or the expectation value of the information gain associated with ascertaining the value of the random variable X . Alternatively, we can think of the entropy as a measure of the amount of uncertainty about X before we ascertain its value. A source that produces one of two distinguishable symbols with equal probability, such as the toss of a fair coin, is said to have a Shannon entropy of 1 bit: ascertaining which symbol is produced, or reducing one’s uncertainty about which symbol is produced, is associated with an amount of information equal to 1 bit.¹ If we already know which symbol will be produced (so the probabilities are 1 and 0), the entropy is 0: there is no uncertainty, and no information gain.

Since

$$(3) \quad p(x_1, \dots, x_n) = 2^{-nH(X)}$$

¹Note that the term ‘bit’ (for ‘binary digit’) is used to refer to the basic unit of classical information in terms of Shannon entropy, and to an elementary two-state classical system considered as representing the possible outputs of an elementary classical information source.

for sufficiently long typical sequences, and the probability of all the typical n -length sequences is less than 1, it follows that there are at most $2^{nH(X)}$ typical sequences. In fact, if the p_i are not all equal, the typical sequences comprise an exponentially small set T (of equiprobable typical sequences) in the set of all sequences as $n \rightarrow \infty$, but since the probability that the source produces an atypical sequence tends to zero as $n \rightarrow \infty$, the set of typical sequences has probability close to 1. So each typical n -sequence could be encoded as a distinct binary number of $nH(X)$ binary digits or bits before being sent through the channel to the receiver, where the original sequence could then be reconstructed by inverting the 1–1 encoding map. (The reconstruction would fail, with low probability, only for the rare atypical sequences, each of which could be encoded as, say, a string of 0's.)

Notice that if the probabilities p_i are all equal ($p_i = 1/k$ for all i), then $H(X) = \log k$, and if some $p_j = 1$ (and so $p_i = 0$ for $i \neq j$), then $H(X) = 0$ (taking $0 \log 0 = \lim_{x \rightarrow 0} x \log x = 0$). It can easily be shown that:

$$(4) \quad 0 \leq H(X) \leq \log k.$$

If we encoded each of the k distinct symbols as a distinct binary number, i.e., as a distinct string of 0's and 1's, we would need binary numbers composed of $\log k$ bits to represent each symbol ($2^{\log k} = k$). So Shannon's analysis shows that messages produced by a stochastic source can be compressed, in the sense that (as $n \rightarrow \infty$ and the probability of an atypical n -length sequence tends to zero) n -length sequences can be encoded without loss of information using $nH(X)$ bits rather than the $n \log k$ bits required if we encoded each of the k symbols a_i as a distinct string of 0's and 1's: this *is* a compression, since $nH(X) < n \log k$ except for equiprobable distributions.

More precisely, let $\bar{X} = \frac{1}{n}(X_1 + X_2 + \dots + X_n)$, where X_1, X_2, \dots, X_n are n independent and identically distributed random variables with mean $\langle X \rangle$ and finite variance. The weak law of large numbers tells us that, for any $\epsilon, \delta > 0$,

$$(5) \quad \Pr(|\bar{X} - \langle X \rangle| \geq \delta) < \epsilon$$

for sufficiently large n .

Now consider a random variable X that takes values x in an alphabet \mathcal{X} with probabilities $p(x) = \Pr(X = x), x \in \mathcal{X}$.² Let

$$(6) \quad Z = -\log p(X)$$

be a function of X that takes the value $-\log p(x)$ when X takes the value x . Then

$$(7) \quad \langle Z \rangle = - \sum_{x \in \mathcal{X}} p(x) \log p(x) = H(X)$$

²Note that $p(x)$ is an abbreviation for $p_X(x)$, so $p(x)$ and $p(y)$ refer to two different random variables. The expression $\Pr(X \in S) = \sum_{x \in S} p(x)$ denotes the probability that the random variable X takes a value in the set S , and $\Pr(X = x)$ denotes the probability that X takes the value x . The expression $p(x_1, x_2, \dots, x_n)$ denotes the probability that the sequence of random variables X_1, X_2, \dots, X_n takes the sequence of values (x_1, x_2, \dots, x_n) . The discussion here follows Cover and Thomas [1991] and I use their notation.

and for a sequence of n independent and identically distributed random variables X_1, X_2, \dots, X_n :

$$\begin{aligned} -\frac{1}{n} \log p(X_1, \dots, X_n) &= -\frac{1}{n} \sum_i \log p(X_i) \\ &= \frac{1}{n} (Z_1 + \dots + Z_n) \\ (8) \qquad \qquad \qquad &= \bar{Z}. \end{aligned}$$

So, by the weak law of large numbers, for $\epsilon, \delta > 0$ and sufficiently large n :

$$(9) \quad \Pr(|\bar{Z} - \langle Z \rangle| \geq \delta) < \epsilon$$

i.e.,

$$(10) \quad \Pr(|-\frac{1}{n} \log p(X_1, \dots, X_n) - H(X)| \geq \delta) < \epsilon$$

or equivalently,

$$(11) \quad \Pr(|-\frac{1}{n} \log p(X_1, \dots, X_n) - H(X)| < \delta) \geq 1 - \epsilon$$

and hence, with probability greater than or equal to $1 - \epsilon$:

$$(12) \quad -n(H(X) + \delta) < \log p(X_1, \dots, X_n) < -n(H(X) - \delta).$$

A ' δ -typical n -length sequence' $(x_1, \dots, x_n) \in \mathcal{X}^n$ of values of the random variables X_1, \dots, X_n is defined as a sequence of symbols of \mathcal{X} satisfying:

$$(13) \quad 2^{-n(H(X)+\delta)} \leq p(x_1, \dots, x_n) \leq 2^{-n(H(X)-\delta)}.$$

Denote the set of δ -typical n -length sequences by $T_\delta^{(n)}$ and the number of sequences in $T_\delta^{(n)}$ by $|T_\delta^{(n)}|$. Then, for sufficiently large n ,

$$(14) \quad \Pr(\{X_1, \dots, X_n\} \in T_\delta^{(n)}) \geq 1 - \epsilon;$$

and it can be shown that

$$(15) \quad (1 - \epsilon)2^{n(H(X)-\delta)} \leq |T_\delta^{(n)}| \leq 2^{n(H(X)+\delta)}.$$

So, roughly, $T^{(n)}$ contains 2^{nH} equiprobable sequences, each having a probability of 2^{-nH} .

Shannon's source coding theorem applies the above result about typical sequences to show that the compression rate of $H(X)$ bits per symbol produced by a source of independent and identically distributed random variables is optimal. The source produces n -length sequences of symbols x_1, x_2, \dots, x_n with probability $p(x_1, x_2, \dots, x_n) = p(x_1)p(x_2)\dots p(x_n)$, where each symbol is chosen from an alphabet \mathcal{X} . If there are k symbols in \mathcal{X} , these n -sequences can be represented as sequences of $n \log k$ bits. Suppose there is a 'block coding' compression scheme

that encodes each ‘block’ or n -length sequence (for sufficiently large n) as a shorter sequence of nR bits, where $0 \leq R \leq \log k$. Suppose also that the receiver has a decompression scheme for decoding sequences of nR bits into sequences of n symbols. Then one speaks of a compression/decompression scheme of *rate* R .

The source coding theorem states that

if the Shannon entropy of a source is $H(X)$, then there exists a reliable compression/decompression scheme of rate R if and only if $R \geq H(X)$, where a scheme is said to be reliable if it reproduces the original sequence with a probability that tends to 1 as $n \rightarrow \infty$.

For reliable communication, we want the compression and decompression of a sequence of symbols to yield the original sequence, but in general there will be a certain probability, $q(x_1, \dots, x_n)$, of decoding a given sequence of nR encoded bits received by the receiver as the original n -sequence produced by the source. The average *fidelity*³ of a compression/decompression scheme for n -length blocks is defined as:

$$(16) \quad F_n = \sum_{\text{all } n\text{-sequences}} p(x_1, \dots, x_n)q(x_1, \dots, x_n)$$

If all the probabilities $q(x_1, \dots, x_n)$ are 1, $F_n = 1$; otherwise $F_n < 1$. In terms of the fidelity as a measure of reliability of correct decoding, the source coding theorem states that

for any $\epsilon, \delta > 0$: (i) there exists a compression/decompression scheme using $H(X) + \delta$ bits per symbol for n -length sequences produced by the source that can be decompressed by the receiver with a fidelity $F_n > 1 - \epsilon$, for sufficiently large n , and (ii) any compression/decompression scheme using $H(X) - \delta$ bits per symbol for n -length sequences will have a fidelity $F_n < \epsilon$, for sufficiently large n .

As a simple example of compression, consider an information source that produces sequences of symbols from a 4-symbol alphabet a_1, a_2, a_3, a_4 with probabilities $1/2, 1/4, 1/8, 1/8$. Each symbol can be represented by a distinct 2-digit binary number:

a_1	:	00
a_2	:	01
a_3	:	10
a_4	:	11

³Note that this definition of fidelity is different from the definition proposed by Nielsen and Chuang [2000, 400] for the fidelity between two probability distributions $\{p_x\}$ and $\{q_x\}$ as a ‘distance measure’ between the distributions. They define $F_{NC}(p_x, q_x) := \sum_x \sqrt{p_x q_x}$, so $F_{NC}(p_x, q_x) = 1$ if $p_x = q_x$.

so without compression we need two bits per symbol of storage space to store the output of the source. The Shannon entropy of the source is $H(X) = -\frac{1}{2} \log \frac{1}{2} - \frac{1}{4} \log \frac{1}{4} - \frac{1}{8} \log \frac{1}{8} - \frac{1}{8} \log \frac{1}{8} = \frac{7}{4}$. Shannon's source coding theorem tells us that there is a compression scheme that uses an average of 7/4 bits per symbol rather than two bits per symbol, and that such a compression scheme is optimal. The optimal scheme is provided by the following encoding:

$$\begin{aligned} a_1 & : 0 \\ a_2 & : 10 \\ a_3 & : 110 \\ a_4 & : 111 \end{aligned}$$

for which the *average* length of a compressed sequence is: $\frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 2 + \frac{1}{8} \cdot 3 + \frac{1}{8} \cdot 3 = \frac{7}{4}$ bits per symbol.

The significance of Shannon's source coding theorem lies in showing that there is an optimal or most efficient way of compressing messages produced by a source (assuming a certain idealization) in such a way that they can be reliably reconstructed by a receiver. Since a message is abstracted as a sequence of distinguishable symbols produced by a stochastic source, the only relevant feature of a message with respect to reliable compression and decompression is the sequence of probabilities associated with the individual symbols: the nature of the physical systems embodying the representation of the message through their states is irrelevant to this notion of compression (provided only that the states are reliably distinguishable), as is the content or meaning of the message. The Shannon entropy $H(X)$ is a measure of the minimal physical resources, in terms of the average number of bits per symbol, that are necessary and sufficient to reliably store the output of a source of messages. In this sense, it is a measure of the amount of information per symbol produced by an information source.

The essential notion underlying Shannon's measure of information is compressibility: information as a physical resource is something that can be compressed, and the amount of information produced by an information source is measured by its optimal compressibility.

2.2 Conditional Entropy, Mutual Information, Channel Capacity

The analysis so far assumes a noiseless channel between the source and the receiver. I turn now to a brief sketch of some concepts relevant to a noisy channel, and a statement of Shannon's noisy channel coding theorem.

An information channel maps inputs consisting of values of a random variable X onto outputs consisting of values of a random variable Y , and the map will generally not be 1-1 if the channel is noisy. Consider the conditional probabilities $p(y|x)$ of obtaining an output value y for a given input value x , for all x, y . From the probabilities $p(x)$ we can calculate $p(y)$ as:

$$p(y) = \sum_x p(y|x)p(x)$$

and we can also calculate $p(x|y)$ by Bayes' rule from the probabilities $p(y|x)$ and $p(x)$, for all x, y , and hence the Shannon entropy of the conditional distribution $p(x|y)$, for all x and a fixed y , denoted by $H(X|Y = y)$.

The quantity

$$(17) \quad H(X|Y) = \sum_y p(y)H(X|Y = y)$$

is known as the *conditional entropy*. It is the expected value of $H(X|Y = y)$ for all y . If we think of $H(X)$, the entropy of the distribution $\{p(x) : x \in \mathcal{X}\}$, as a measure of the uncertainty of the X -value, then $H(X|Y = y)$ is a measure of the uncertainty of the X -value, given the Y -value y , and $H(X|Y)$ is a measure of the average uncertainty of the X -value, given a Y -value.

Putting it differently, the number of input sequences of length n that are consistent with a given output sequence (as $n \rightarrow \infty$) is $2^{nH(X|Y)}$, i.e., $H(X|Y)$ is the number of bits per symbol of additional information needed, on average, to identify an input X -sequence from a given Y -sequence. This follows because there are $2^{nH(X,Y)}$ typical sequences of pairs (x, y) , where the *joint entropy* $H(X, Y)$ is calculated from the joint probability $p(x, y)$. So there are

$$(18) \quad \frac{2^{nH(X,Y)}}{2^{nH(Y)}} = 2^{n(H(X,Y)-H(Y))} = 2^{nH(X|Y)}$$

typical X -sequences associated with a given Y -sequence.

The 'chain rule' equality

$$(19) \quad H(X, Y) = H(X) + H(Y|X) = H(Y) + H(X|Y) = H(Y, X)$$

follows immediately from the logarithmic definitions of the quantities:

$$\begin{aligned} H(X:Y) &:= - \sum_{x,y} p(x, y) \log p(x, y) \\ &= - \sum_{x,y} p(x)p(y|x) \log (p(x)p(y|x)) \\ &= - \sum_{x,y} p(x)p(y|x) \log p(x) - \sum_{x,y} p(x)p(y|x) \log p(y|x) \\ &= - \sum_x p(x) \log p(x) + \sum_x p(x) \left(- \sum_y p(y|x) \log p(y|x) \right) \\ (20) \quad &= H(X) + H(Y|X) \end{aligned}$$

Note that $H(X|Y) \neq H(Y|X)$.

The *mutual information* measures the average amount of information gained about X by ascertaining a Y -value, i.e., the amount of information one random variable contains about another, or the reduction in uncertainty of one random variable obtained by measuring another.

Mutual information can be defined in terms of the concept of *relative entropy*, which is a measure of something like the distance between two probability distributions (although it is not a true metric, since it is not symmetric and does not satisfy the triangle inequality). The relative entropy between distributions $p(x)$ and $q(x)$ is defined as:

$$(21) \quad D(p \parallel q) = \sum_{x \in \mathcal{A}} p(x) \log \frac{p(x)}{q(x)}.$$

The mutual information can now be defined as:

$$(22) \quad \begin{aligned} H(X:Y) &= D(p(x,y) \parallel p(x)p(y)) \\ &= \sum_x \sum_y p(x,y) \log \frac{p(x,y)}{p(x)p(y)}. \end{aligned}$$

It follows that

$$(23) \quad H(X:Y) = H(X) - H(X|Y) = H(Y) - H(Y|X),$$

i.e., the mutual information of two random variables represents the average information gain about one random variable obtained by measuring the other: the difference between the initial uncertainty of one of the random variables, and the average residual uncertainty of that random variable after ascertaining the value of the other random variable. Also, since $H(X,Y) = H(X) + H(Y|X)$, it follows that

$$(24) \quad H(X:Y) = H(X) + H(Y) - H(X,Y);$$

i.e., the mutual information of two random variables is a measure of how much information they have in common: the sum of the information content of the two random variables, as measured by the Shannon entropy (in which joint information is counted twice), minus their joint information. Note that $H(X:X) = H(X)$, as we would expect.

For a noisy channel, if X represents the input to the channel and Y represents the output of the channel, $H(X:Y)$ represents the average amount of information gained about the input X by ascertaining the value of the output Y . The *capacity* of a channel, C , is defined as the supremum of $H(X:Y)$ over all input distributions.

Shannon's noisy channel coding theorem shows, perhaps surprisingly, that up to C bits of information can be sent through a noisy channel with arbitrary low error rate. That is,

there exists an optimal coding for an information source with entropy $H \leq C$ such that n -length sequences produced by the source can be transmitted faithfully over the channel: the error rate tends to zero as $n \rightarrow \infty$. The probability of error tends to 1 if we attempt to transmit more than C bits through the channel.

This means that there are two ways of improving the transmission rate over a noisy channel such as a telephone cable. We can improve the channel capacity by replacing the cable with a faster one, or we can improve the information processing (the data compression).

3 QUANTUM INFORMATION

The physical notion of information, discussed in §2, is profoundly transformed by the transition from classical mechanics to quantum mechanics. The aim of this section is to bring out the nature of this transformation. In §3.1, I develop some core concepts of quantum mechanics relevant to quantum information: entangled states, the Schmidt decomposition, the density operator formalism for the representation of pure and mixed states, the ‘purification’ of mixed states, generalized quantum measurements in terms of positive operator valued measures (POVMs), and the evolution of open systems represented by quantum operations. I assume throughout Hilbert spaces of finite dimension (and so avoid all the technicalities of functional analysis required for the treatment of infinite-dimensional Hilbert spaces). In fact, there is no loss of generality here, since both classical and quantum information sources are considered to produce messages consisting of sequences of symbols from some finite alphabet, which we represent in terms of a finite set of classical or quantum states. Moreover, all the conceptual issues relevant to the difference between classical and quantum information show up in finite-dimensional Hilbert spaces. In §3.2, I introduce von Neumann’s generalization of the Shannon entropy and related notions for quantum information. In §3.3 and §3.4, I discuss some salient features that distinguish quantum information from classical information: §3.3 deals with the limitations on copying quantum information imposed by the ‘no cloning’ theorem, and §3.4 deals with the limited accessibility of quantum information defined by the Holevo bound. Finally, in §3.5 I show how the notion of compressibility applies to quantum information, and I outline Schumacher’s generalization of Shannon’s source coding theorem for quantum information, noting a distinction between ‘visible’ and ‘blind’ compression applicable to quantum information.

3.1 *Some Relevant Quantum Mechanics*

Entangled States

Consider a quantum system Q which is part of a compound system QE ; E for ‘environment’, although E could be any quantum system of which Q is a subsystem. Pure states of QE are represented as rays or unit vectors in a tensor product Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^E$. A general pure state of QE is a state of the form:

$$(25) \quad |\Psi\rangle = \sum c_{ij} |q_i\rangle |e_j\rangle$$

where $|q_i\rangle \in \mathcal{H}^Q$ is a complete set of orthonormal states (a basis) in \mathcal{H}^Q and $|e_j\rangle \in \mathcal{H}^E$ is a basis in \mathcal{H}^E . If the coefficients c_{ij} are such that $|\Psi\rangle$ cannot be expressed as a product state $|Q\rangle|E\rangle$, then $|\Psi\rangle$ is called an *entangled state*.

For any state $|\Psi\rangle$ of QE , there exist orthonormal bases $|i\rangle \in \mathcal{H}^Q$, $|j\rangle \in \mathcal{H}^E$ such that $|\Psi\rangle$ can be expressed in a biorthogonal correlated form as:

$$(26) \quad |\Psi\rangle = \sum_i \sqrt{p_i} |i\rangle |i\rangle$$

where the coefficients $\sqrt{p_i}$ are real and non-negative, and $\sum p_i = 1$. This representation is referred to as the *Schmidt decomposition*. The Schmidt decomposition is unique if and only if the p_i are all distinct.

An example is the biorthogonal EPR state⁴

$$(27) \quad |\Psi\rangle = (|0\rangle|1\rangle - |1\rangle|0\rangle)/\sqrt{2};$$

say, the singlet state of two spin-1/2 particles (the Schmidt form with positive coefficients is obtained by absorbing the relative phases in the definition of the basis vectors). In the singlet state, $|0\rangle$ and $|1\rangle$ can be taken as representing the two eigenstates of spin in the z -direction, but since the state is symmetric, $|\Psi\rangle$ retains the same form for spin in any direction. The EPR argument exploits the fact that spin measurements in the same direction on the two particles, which could be arbitrarily far apart, will yield outcomes that are perfectly anti-correlated for any spin direction. Bell's counterargument exploits the fact that when the spin is measured on one particle in a direction θ_1 to the z -axis, but on the other particle in a direction θ_2 to the z -axis, the probability of finding the same outcome for both particles (both 1 or both 0) is $\frac{1}{2} \sin^2(\theta_1 - \theta_2)$. It follows that the outcomes are perfectly correlated when $\theta_1 - \theta_2 = \pi$ and that 3/4 of the outcomes are the same when $\theta_1 - \theta_2 = 2\pi/3$. On the other hand, from Bell's inequality, derived under Einstein's realist assumptions of separability and locality, we see that the correlation for $\theta_1 - \theta_2 = 2\pi/3$ cannot exceed 2/3. See Dickson (this vol., ch. 4) for further discussion.

This means that the dynamical evolution of a quantum system can result in a state representing correlational information that no classical computer can simulate. That is, no classical computer can be programmed to perform the following task: for any pair of input angles, θ_1, θ_2 , at different locations, output a pair of values (0 or 1) for these locations such that the values are perfectly correlated when $\theta_1 - \theta_2 = \pi$, perfectly anti-correlated when $\theta_1 = \theta_2$, and 75% correlated when $\theta_1 - \theta_2 = 2\pi/3$, where the response time between being given the input and producing the output in each case is less than the time taken by light to travel between the two locations.

⁴Einstein, Podolsky and Rosen considered a more complicated state entangled over position and momentum values. The spin example is due to Bohm [1951, pp, 611–623].

Notice that the four states:

$$(28) \quad |00\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle)$$

$$(29) \quad |01\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|0\rangle)$$

$$(30) \quad |10\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle - |1\rangle|1\rangle)$$

$$(31) \quad |11\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)$$

form an orthonormal basis, called the Bell basis, in the 2×2 -dimensional Hilbert space. Any Bell state can be transformed into any other Bell state by a local unitary transformation, X, Y , or Z , where X, Y, Z are the Pauli spin matrices:

$$(32) \quad X = \sigma_x = |0\rangle\langle 1| + |1\rangle\langle 0| = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

$$(33) \quad Y = \sigma_y = i|0\rangle\langle 1| - i|1\rangle\langle 0| = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$$

$$(34) \quad Z = \sigma_z = |0\rangle\langle 0| + |1\rangle\langle 1| = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

For example:

$$(35) \quad X \otimes I \cdot \frac{1}{\sqrt{2}}(|0\rangle\langle 1| - |1\rangle\langle 0|) = \frac{1}{\sqrt{2}}(|0\rangle\langle 0| - |1\rangle\langle 1|).$$

If QE is a closed system in an entangled pure state represented by

$$(36) \quad |\Psi\rangle = \sum_i \sqrt{p_i} |i\rangle |i\rangle$$

in the Schmidt decomposition, the expected value of any Q -observable A on \mathcal{H}^Q can be computed as:

$$\begin{aligned} \langle A \rangle &= \text{Tr}(|\Psi\rangle\langle\Psi|A \otimes I) \\ &= \text{Tr}_Q(\text{Tr}_E(|\Psi\rangle\langle\Psi|A)) \\ &= \text{Tr}_Q\left(\sum_i p_i |i\rangle\langle i|A\right) \\ (37) \quad &= \text{Tr}_Q(\rho A) \end{aligned}$$

where $\text{Tr}_Q() = \sum_q \langle q_i | \cdot | q_i \rangle$, for any orthonormal basis in \mathcal{H}^Q , is the partial trace over \mathcal{H}^Q , $\text{Tr}_E()$ is the partial trace over \mathcal{H}^E , and $\rho = \sum_i p_i |i\rangle\langle i| \in \mathcal{H}^Q$ is the *reduced density operator* of the open system Q , a positive operator with unit trace.

Since the density operator ρ yields the statistics of all Q -observables via Eq. (37), ρ is taken as representing the quantum state of the system Q .

If QE is an entangled pure state, then the open system Q is in a *mixed state* ρ , i.e., $\rho \neq \rho^2$; for pure states, ρ is a projection operator onto a ray and $\rho = \rho^2$. A mixed state represented by a density operator $\rho = \sum \rho_i |i\rangle\langle i|$ can be regarded as a mixture of pure states $|i\rangle$ prepared with prior probabilities p_i , but this representation is not unique — not even if the states combined in the mixture are orthogonal. For example, the equal-weight mixture of orthonormal states $|0\rangle, |1\rangle$ in a 2-dimensional Hilbert space \mathcal{H}_2 has precisely the same statistical properties, and hence the same density operator $\rho = I/2$, as the equal weight mixture of any pair of orthonormal states, e.g., the states $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, or the equal-weight mixture of nonorthogonal states $|0\rangle, \frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}|1\rangle, \frac{1}{2}|0\rangle - \frac{\sqrt{3}}{2}|1\rangle$ 120° degrees apart, or the uniform continuous distribution over all possible states in \mathcal{H}_2 .

More generally, for any basis of orthonormal states $|e_i\rangle \in \mathcal{H}^E$, the entangled state $|\Psi\rangle$ can be expressed as:

$$(38) \quad |\Psi\rangle = \sum_{ij} c_{ij} |q_i\rangle |e_j\rangle = \sum_j \sqrt{w_j} |r_j\rangle |e_j\rangle$$

where the normalized states $|r_j\rangle = \sum_i \frac{c_{ij}}{\sqrt{w_j}} |q_i\rangle$ are *relative states* to the $|e_j\rangle$ ($\sqrt{w_j} = \sum_j |c_{ij}|^2$). Note that the states $|r_j\rangle$ are not in general orthogonal. Since the $|e_j\rangle$ are orthogonal, we can express the density operator representing the state of Q as:

$$(39) \quad \rho = \sum_i w_i |r_i\rangle\langle r_i|.$$

In effect, a measurement of an E -observable with eigenstates $|e_i\rangle$ will leave the composite system QE in one of the states $|r_i\rangle |e_i\rangle$ with probability w_i , and a measurement of an E -observable with eigenstates $|i\rangle$ (the orthogonal states of the Schmidt decomposition in (36) above) will leave the system QE in one of the states $|i\rangle |i\rangle$ with probability p_i . Since Q and E could be widely separated from each other in space, no measurement at E could affect the statistics of any Q -observable; or else measurements at E would allow superluminal signaling between Q and E . It follows that the mixed state ρ can be realized as a mixture of orthogonal states $|i\rangle$ (the eigenstates of ρ) with weights p_i , or as a mixture of non-orthogonal relative states $|r_j\rangle$ with weights w_j in infinitely many ways, depending on the choice of basis in \mathcal{H}^E :

$$(40) \quad \rho = \sum_i p_i |i\rangle\langle i| = \sum_j w_j |r_j\rangle\langle r_j|$$

and all these different mixtures with the same density operator ρ must be physically indistinguishable.

Note that any mixed state density operator $\rho \in \mathcal{H}^Q$ can be ‘purified’ by adding a suitable ancilla system E , in the sense that ρ is the partial trace of a pure state $|\Psi\rangle \in \mathcal{H}^Q \otimes \mathcal{H}^E$ over \mathcal{H}^E . A purification of a mixed state is, clearly, not unique, but depends on the choice of $|\Psi\rangle$ in \mathcal{H}^E . The Hughston-Jozsa-Wootters theorem [Hughston *et al.*, 1993] shows that for *any* mixture of pure states $|r_i\rangle$ with weights w_i , where $\rho = \sum_j w_j |r_j\rangle\langle r_j|$, there is a purification of ρ and a suitable measurement on the system E that will leave Q in the mixture ρ . So an observer at E can remotely prepare Q in any mixture that corresponds to the density operator ρ (and of course all these different mixtures are physically indistinguishable). Similar results were proved earlier by Schrödinger [1936], Jaynes [1957] and Gisin [1989]. See Halvorson [2004] for a generalization to hyperfinite von Neuman algebras.

Measurement

A standard von Neumann ‘yes-no’ measurement is associated with a projection operator; so a standard observable is represented in the spectral representation as a sum of projection operators, with coefficients representing the eigenvalues of the observable. Such a measurement is the quantum analogue of the measurement of a property of a system in classical physics. Classically, we think of a property of a system as being associated with a subset in the state space (phase space) of the system, and determining whether the system has the property amounts to determining whether the state of the system lies in the corresponding subset. In quantum mechanics, the counterpart of a subset in phase space is a closed linear subspace in Hilbert space. Just as the different possible values of an observable (dynamical quantity) of a classical system correspond to the subsets in a mutually exclusive and collectively exhaustive set of subsets covering the classical state space, so the different values of a quantum observable correspond to the subspaces in a mutually exclusive (i.e., orthogonal) and collectively exhaustive set of subspaces spanning the quantum state space. (For further discussion, see Dickson, this vol., ch. 4, [Mackey, 1963], and [Bub, 1997].)

In quantum mechanics, and especially in the theory of quantum information (where any read-out of the quantum information encoded in a quantum state requires a quantum measurement), it is useful to consider a more general class of measurements than the projective measurements associated with the determination of the value of an observable. It is common to speak of generalized measurements and generalized observables. But in fact this terminology is more misleading than illuminating, because a generalized measurement is not a procedure that reveals whether or not a quantum system has some sort of generalized property. Rather, the point of the generalization is to exploit the difference between quantum and classical states for new possibilities in the representation and manipulation of information.

To clarify the idea, I will follow the excellent discussion by Nielsen and Chuang [2000, §2.2.3–2.2.6]. A quantum measurement can be characterized, completely generally, as a certain sort of interaction between two quantum systems, Q (the

measured system) and M (the measuring system). We suppose that Q is initially in a state $|\psi\rangle$ and that M is initially in some standard state $|0\rangle$, where $|m\rangle$ is an orthonormal basis of ‘pointer’ eigenstates in \mathcal{H}^M . The interaction is defined by a unitary transformation U on the Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^M$ that yields the transition:

$$(41) \quad |\psi\rangle|0\rangle \xrightarrow{U} \sum_m M_m|\psi\rangle|m\rangle$$

where $\{M_m\}$ is a set of linear operators (the Kraus operators) defined on \mathcal{H}^Q satisfying the *completeness condition*:

$$(42) \quad \sum_m M_m^\dagger M_m = I.$$

(The symbol \dagger denotes the adjoint or Hermitian conjugate.) The completeness condition guarantees that this evolution is unitary, because it guarantees that U preserves inner products, i.e.

$$\begin{aligned} \langle\phi|\langle 0|U^\dagger U|\psi\rangle|0\rangle &= \sum_{m,m'} \langle m|\langle\phi|M_m^\dagger M_{m'}|\psi\rangle|m'\rangle \\ &= \sum_m \langle\phi|M_m^\dagger M|\psi\rangle \\ (43) \qquad \qquad \qquad &= \langle\phi|\psi\rangle \end{aligned}$$

from which it follows that U , defined as above by Eq. (41) for any product state $|\psi\rangle|0\rangle$ (for any $|\psi\rangle \in \mathcal{H}^Q$) can be extended to a unitary operator on the Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^M$. Accordingly, any set of linear operators $\{M_m\}$ defined on the Hilbert space of the system Q satisfying the completeness condition defines a measurement in this general sense, with the index m labeling the possible outcomes of the measurement, and any such set is referred to as a set of measurement operators.

If we now perform a standard projective measurement on M to determine the value m of the pointer observable, defined by the projection operator

$$P_m = I_Q \otimes |m\rangle\langle m|$$

then the probability of obtaining the outcome m is, by (37)⁵:

$$\begin{aligned} p(m) &= \langle 0|\langle\psi|U^\dagger P_m U|\psi\rangle|0\rangle \\ &= \sum_{m'm''} \langle m'|\langle\psi|M_{m'}^\dagger (I_Q \otimes |m\rangle\langle m|) M_{m''}|\psi\rangle|m''\rangle \\ &= \sum_{m'm''} \langle\psi|M_{m'}^\dagger \langle m'|m\rangle\langle m|m''\rangle M_{m''}|\psi\rangle \\ (44) \qquad \qquad \qquad &= \langle\psi|M_m^\dagger M_m|\psi\rangle; \end{aligned}$$

⁵The expected value of a projection operator, which is an idempotent observable with eigenvalues 0 and 1, is equal to the probability of obtaining the eigenvalue 1. Here the eigenvalue 1 corresponds to the outcome m .

and, more generally, if the initial state of Q is a mixed state ρ , then

$$(45) \quad p(m) = \text{Tr}_Q(M\rho M^\dagger).$$

The final state of QM after the projective measurement on M yielding the outcome m is:

$$(46) \quad \frac{P_m U |\psi\rangle |0\rangle}{\sqrt{\langle \psi | U^\dagger P U | \psi \rangle}} = \frac{M_m |\psi\rangle |m\rangle}{\sqrt{\langle \psi | M_m^\dagger M_m | \psi \rangle}}.$$

So the final state of M is $|m\rangle$ and the final state of Q is:

$$\frac{M_m |\psi\rangle}{\sqrt{\langle \psi | M_m^\dagger M_m | \psi \rangle}};$$

and, more generally, if the initial state of Q is a mixed state ρ , then the final state of Q is:

$$\frac{M_m \rho M_m^\dagger}{\text{Tr}_Q(M_m \rho M_m^\dagger)}.$$

Note that this general notion of measurement covers the case of standard projective measurements. In this case $\{M_m\} = \{P_m\}$, where $\{P_m\}$ is the set of projection operators defined by the spectral measure of a standard quantum observable represented by a self-adjoint operator. It also covers the measurement of ‘generalized observables’ associated with positive operator valued measures (POVMs). Let

$$(47) \quad E_m = M_m^\dagger M_m$$

then the set $\{E_m\}$ defines a set of positive operators (‘effects’) such that

$$(48) \quad \sum E_m = I$$

A POVM can be regarded as a generalization of a projection valued measure (PVM), in the sense that Eq. (48) defines a ‘resolution of the identity’ without requiring the PVM orthogonality condition:

$$(49) \quad P_m P_{m'} = \delta_{mm'} P_m.$$

Note that for a POVM:

$$(50) \quad p(m) = \langle \psi | E_m | \psi \rangle.$$

Given a set of positive operators $\{E_m\}$ such that $\sum E_m = I$, measurement operators M_m can be defined via

$$(51) \quad M_m = U \sqrt{E_m},$$

where U is a unitary operator, from which it follows that

$$(52) \quad \sum_m M_m^\dagger M_m = \sum E_m = I$$

As a special case, of course, we can take $U = 1$ and $M_m = \sqrt{E_m}$. Conversely, given a set of measurement operators $\{M_m\}$, there exist unitary operators U_m such that $M_m = U_m \sqrt{E_m}$, where $\{E_m\}$ is a POVM. (This follows immediately from [Nielsen and Chuang, 2000, Theorem 2.3, p. 78]; see [Nielsen and Chuang, 2000, Exercise 2.63, p. 92].)

Except for the standard case of projective measurements, one might wonder why it might be useful to single out such unitary transformations, and why in the general case such a process should be called a *measurement* of Q . The following example, taken from [Nielsen and Chuang, 2000, 92], is illuminating. Suppose we know that a system with a 2-dimensional Hilbert space is in one of two nonorthogonal states:

$$\begin{aligned} |\psi_1\rangle &= |0\rangle \\ |\psi_2\rangle &= \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \end{aligned}$$

It is impossible to *reliably* distinguish these states by a quantum measurement, even in the above generalized sense. Here ‘reliably’ means that the state is identified correctly with zero probability of error.

To see this, suppose there is such a measurement, defined by two measurement operators M_1, M_2 satisfying the completeness condition. Then we require

$$(53) \quad p(1) = \langle \psi_1 | M_1^\dagger M_1 | \psi_1 \rangle = 1,$$

to represent reliability if the state is $|\psi_1\rangle$; and

$$(54) \quad p(2) = \langle \psi_2 | M_2^\dagger M_2 | \psi_2 \rangle = 1$$

to represent reliability if the state is $|\psi_2\rangle$. By the completeness condition we must have

$$(55) \quad \langle \psi_1 | M_1^\dagger M_1 + M_2^\dagger M_2 | \psi_1 \rangle = 1$$

from which it follows that $\langle \psi_1 | M_2^\dagger M_2 | \psi_1 \rangle = 0$, i.e., $M_2 |\psi_1\rangle = M_2 |0\rangle = 0$. Hence

$$(56) \quad M_2 |\psi_2\rangle = M_2 \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = \frac{1}{\sqrt{2}} M_2 |1\rangle$$

and so

$$(57) \quad p(2) = \langle \psi_2 | M_2^\dagger M_2 | \psi_2 \rangle = \frac{1}{2} \langle 1 | M_2^\dagger M_2 | 1 \rangle.$$

But by the completeness condition we also have

$$(58) \quad \langle 1 | M_2^\dagger M_2 | 1 \rangle \leq \langle 1 | M_1^\dagger M_1 + M_2^\dagger M_2 | 1 \rangle = \langle 1 | 1 \rangle = 1$$

from which it follows that

$$(59) \quad p(2) \leq \frac{1}{2}$$

which contradicts Eq. (54).

However, it is possible to perform a measurement in the generalized sense, with *three* possible outcomes, that will allow us to correctly identify the state some of the time, i.e., for two of the possible outcomes, while nothing about the identity of the state can be inferred from the third outcome.

Here's how: The three operators

$$(60) \quad \begin{aligned} E_1 &= \frac{\sqrt{2}}{1+\sqrt{2}} \frac{(|0\rangle - |1\rangle)(\langle 0| - \langle 1|)}{2} \\ E_2 &= \frac{\sqrt{2}}{1+\sqrt{2}} |1\rangle\langle 1| \\ E_3 &= I - E_1 - E_2 \end{aligned}$$

are all positive operators and $E_1 + E_2 + E_3 = I$, so they define a POVM. In fact, E_1, E_2, E_3 are each multiples of projection operators onto the states

$$(61) \quad \begin{aligned} |\phi_1\rangle &= |\psi_2\rangle^\perp \\ |\phi_2\rangle &= |\psi_1\rangle^\perp \\ |\phi_3\rangle &= \frac{(1+\sqrt{2})|0\rangle + |1\rangle}{\sqrt{2\sqrt{2}(1+\sqrt{2})}} \end{aligned}$$

with coefficients $\frac{\sqrt{2}}{1+\sqrt{2}}, \frac{\sqrt{2}}{1+\sqrt{2}}, \frac{1}{1+\sqrt{2}}$ respectively. The measurement involves a system M with three orthogonal pointer states $|1\rangle, |2\rangle, |3\rangle$. The appropriate unitary interaction U results in the transition, for an input state $|\psi\rangle$:

$$(62) \quad |\psi\rangle|0\rangle \xrightarrow{U} \sum_m M_m |\psi\rangle|m\rangle$$

where $M_m = \sqrt{E_m}$.

If the input state is $|\psi_1\rangle = |0\rangle$, we have the transition:

$$(63) \quad \begin{aligned} |\psi_1\rangle|0\rangle &\xrightarrow{U} \sqrt{E_1}|0\rangle|1\rangle + \sqrt{E_3}|0\rangle|3\rangle \\ &= \alpha|\phi_1\rangle|1\rangle + \beta|\phi_3\rangle|3\rangle \end{aligned}$$

(because $\sqrt{E_2}|\psi_1\rangle = \sqrt{E_2}|0\rangle = 0$). And if the input state is $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, we have the transition:

$$(64) \quad \begin{aligned} |\psi_2\rangle|0\rangle &\xrightarrow{U} \sqrt{E_2} \frac{|0\rangle + |1\rangle}{\sqrt{2}} |2\rangle + \sqrt{E_3} \frac{|0\rangle + |1\rangle}{\sqrt{2}} |3\rangle \\ &= \gamma|\phi_2\rangle|2\rangle + \delta|\phi_3\rangle|3\rangle \end{aligned}$$

(because $\sqrt{E_1}|\psi_2\rangle = \sqrt{E_1}\frac{|0\rangle+|1\rangle}{\sqrt{2}} = 0$), where $\alpha, \beta, \gamma, \delta$ are real numerical coefficients.

We see that a projective measurement of the pointer of M that yields the outcome $m = 1$ indicates, with certainty, that the input state was $|\psi_1\rangle = |0\rangle$. In this case, the measurement leaves the system Q in the state $|\phi_1\rangle$. A measurement outcome $m = 2$ indicates, with certainty, that the input state was $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and in this case the measurement leaves the system Q in the state $|\phi_2\rangle$. If the outcome is $m = 3$, the input state could have been either $|\psi_1\rangle = |0\rangle$ or $|\psi_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, and Q is left in the state $|\phi_3\rangle$.

Quantum Operations

When a closed system QE evolves under a unitary transformation, Q can be shown to evolve under a *quantum operation*, i.e., a completely positive linear map:

$$(65) \quad \mathcal{E} : \rho \rightarrow \rho'$$

where

$$(66) \quad \mathcal{E}(\rho) = \text{Tr}_E(U\rho \otimes \rho_E U^\dagger)$$

(See [Nielsen and Chuang, 2000, 356 ff].) The map \mathcal{E} is linear (or convex-linear) in the sense that $\mathcal{E}(\sum_i p_i \rho_i) = \sum_i p_i \mathcal{E}(\rho_i)$, positive in the sense that \mathcal{E} maps positive operators to positive operators, and completely positive in the sense that $\mathcal{E} \otimes I$ is a positive map on the extension of \mathcal{H}^Q to a Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^E$, associated with the addition of any ancilla system E to Q .

Every quantum operation (i.e., completely positive linear map) on a Hilbert space \mathcal{H}^Q has a (non-unique) representation as a unitary evolution on an extended Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^E$, i.e.,

$$(67) \quad \mathcal{E}(\rho) = \text{Tr}_E(U(\rho \otimes \rho_E)U^\dagger)$$

where ρ_E is an appropriately chosen initial state of an ancilla system E (which we can think of as the environment of Q). It turns out that it suffices to take ρ_E as a pure state, i.e., $|0\rangle\langle 0|$, since a mixed state of E can always be purified by enlarging the Hilbert space (i.e., adding a further ancilla system). So the evolution of a system Q described by a quantum operation can always be modeled as the unitary evolution of a system QE , for an initial pure state of E .

Also, every quantum operation on a Hilbert space \mathcal{H}^Q has a (non-unique) operator sum representation intrinsic to \mathcal{H}^Q :

$$(68) \quad \mathcal{E}(\rho) = \sum_i E_i \rho E_i^\dagger$$

where $E_i = \langle i|U|0\rangle$ for some orthonormal basis $\{|i\rangle\}$ of E . (See [Nielsen and Chuang, 2000, Theorem 8.1, p. 368].) If the operation is trace-preserving (or nonselective), then $\sum_i E_i^\dagger E_i = I$. For operations that are not trace-preserving (or

selective), $\sum_i E_i^\dagger E_i \leq I$. This corresponds to the case where the outcome of a measurement on QE is taken into account (selected) in the transition $\mathcal{E} \rightarrow \mathcal{E}(\rho)$.

If there is no interaction between Q and E , then $\epsilon(\rho) = U_Q \rho U_Q^\dagger$, $U_Q U_Q^\dagger = I$, i.e., there is only one operator in the sum. In this case, $U = U_Q \otimes U_E$ and

$$(69) \quad \mathcal{E}(\rho) = \text{Tr}_E(U_Q \otimes U_E(\rho \otimes |0\rangle\langle 0|)U_Q^\dagger \otimes U_E^\dagger)$$

$$(70) \quad = U_Q \rho U_Q^\dagger.$$

So unitary evolution is a special case of the operator sum representation of a quantum operation and, of course, another special case is the transition $\mathcal{E} \rightarrow \mathcal{E}(\rho)$ that occurs in a quantum measurement process, where $E_i = M_i$. A trace-preserving operation corresponds to a non-selective measurement:

$$(71) \quad \mathcal{E}(\rho) = \sum_i M_i \rho M_i^\dagger;$$

while an operation that is not trace-preserving corresponds to a selective measurement, where the state ‘collapses’ onto the corresponding measurement outcome:

$$(72) \quad M_i \rho M_i^\dagger / \text{Tr}(M_i \rho M_i^\dagger).$$

The operator sum representation applies to quantum operations between possibly different input and output Hilbert spaces, and characterizes the following general situation: a quantum system in an unknown initial state ρ is allowed to interact unitarily with other systems prepared in standard states, after which some part of the composite system is discarded, leaving the final system in a state ρ' . The transition $\rho \rightarrow \rho'$ is defined by a quantum operation. So a quantum operation represents, quite generally, the unitary evolution of a closed quantum system, the nonunitary evolution of an open quantum system in interaction with its environment, and evolutions that result from a combination of unitary interactions and selective or nonselective measurements.

As we have seen, the creed of the Church of the Larger Hilbert Space is that every state can be made pure, every measurement can be made ideal, and every evolution can be made unitary – on a larger Hilbert space.⁶

3.2 Von Neumann Entropy

In this section, I define the von Neumann entropy of a mixture of quantum states (von Neumann’s generalization of the Shannon entropy of a classical probability distribution characterizing a classical information source) and the corresponding notions of conditional entropy and mutual information.

Information in Shannon’s sense is a quantifiable resource associated with the output of a (suitably idealized) stochastic source of symbolic states, where the physical nature of the systems embodying these states is irrelevant to the amount

⁶The Creed originates with John Smolin. I owe this formulation to Ben Schumacher. See his *Lecture Notes on Quantum Information Theory* [1998].

of classical information associated with the source. The quantity of information associated with a stochastic source is defined by its optimal compressibility, and this is given by the Shannon entropy. The fact that some feature of the output of a stochastic source can be optimally compressed is, ultimately, what justifies the attribution of a quantifiable resource to the source.

Information is represented physically in the states of physical systems. The essential difference between classical and quantum information arises because of the different distinguishability properties of classical and quantum states. As we will see below, only sets of orthogonal quantum states are reliably distinguishable (i.e., with zero probability of error), as are sets of different classical states (which are represented by disjoint singleton subsets in a phase space, and so are orthogonal as subsets of phase space in a sense analogous to orthogonal subspaces of a Hilbert space).

Classical information is that sort of information represented in a set of distinguishable states — states of classical systems, or orthogonal quantum states — and so can be regarded as a subcategory of quantum information, where the states may or may not be distinguishable. The idea behind quantum information is to extend Shannon's notion of compressibility to a stochastic source of quantum states, which may or may not be distinguishable. For this we need to define a suitable measure of information for probability distributions of quantum states — mixtures — as a generalization of the notion of Shannon entropy.

Consider a system QE in an entangled state $|\Psi\rangle$. Then the subsystem Q is in a mixed state ρ , which can always be expressed as:

$$(73) \quad \rho = \sum_i p_i |i\rangle\langle i|$$

where the p_i are the eigenvalues of ρ and the pure states $|i\rangle$ are orthonormal eigenstates of ρ . This is the spectral representation of ρ , and any density operator — a positive (hence Hermitian) operator — can be expressed in this way. The representation is unique if and only if the p_i are all distinct. If some of the p_i are equal, there is a unique representation of ρ as a sum of projection operators with the distinct values of the p_i as coefficients, but some of the projection operators will project onto multi-dimensional subspaces.

Since ρ has unit trace, $\sum p_i = 1$, and so the spectral representation of ρ represents a classical probability distribution of orthogonal, and hence distinguishable, pure states. If we measure a Q -observable with eigenstates $|i\rangle$, then the outcomes can be associated with the values of a random variable X , where $\Pr(X = i) = p_i$. Then

$$(74) \quad H(X) = - \sum p_i \log p_i$$

is the Shannon entropy of the probability distribution of measurement outcomes.

Now,

$$(75) \quad -\text{Tr}(\rho \log \rho) = - \sum p_i \log p_i$$

(because the eigenvalues of $\rho \log \rho$ are $p_i \log p_i$ and the trace of an operator is the sum of the eigenvalues), so a natural generalization of Shannon entropy for any mixture of quantum states with density operator ρ is the *von Neumann entropy*⁷:

$$(76) \quad S := -\text{Tr}(\rho \log \rho)$$

which coincides with the Shannon entropy for measurements in the eigenbasis of ρ . For a completely mixed state $\rho = I/d$, where $\dim \mathcal{H}^Q = d$, the d eigenvalues of ρ are all equal to $1/d$ and $S = \log d$. This is the maximum value of S in a d -dimensional Hilbert space. The von Neumann entropy S is zero, the minimum value, if and only if ρ is a pure state, where the eigenvalues of ρ are 1 and 0. So $0 \leq S \leq \log d$, where d is the dimension of \mathcal{H}^Q .

Recall that we can think of the Shannon entropy as a measure of the average amount of information gained by identifying the state produced by a known stochastic source. Alternatively, the Shannon entropy represents the optimal compressibility of the information produced by an information source. The von Neumann entropy does *not*, in general, represent the amount of information gained by identifying the quantum state produced by a stochastic source characterized as a mixed state, because nonorthogonal quantum states in a mixture cannot be reliably identified. However, as we will see in §3.5, the von Neumann entropy can be interpreted in terms of compressibility via Schumacher's source coding theorem Schumacher's source coding theorem for quantum information [Schumacher, 1995], a generalization of Shannon's source coding theorem for classical information. For an elementary two-state quantum system with a 2-dimensional Hilbert space considered as representing the output of an elementary quantum information source, $S = 1$ for an equal weight distribution over two orthogonal states (i.e., for the density operator $\rho = I/2$), so Schumacher takes the basic unit of quantum information as the 'qubit.' By analogy with the term 'bit', the term 'qubit' refers to the basic unit of quantum information in terms of the von Neumann entropy, and to an elementary two-state quantum system considered as representing the possible outputs of an elementary quantum information source.

The difference between quantum information as measured by von Neumann entropy S and classical information as measured by Shannon entropy H can be brought out by considering the quantum notions of conditional entropy and mutual information (cf. §2.2), and in particular the peculiar feature of *inaccessibility* associated with quantum information.

For a composite system AB , conditional von Neumann entropy and mutual information are defined in terms of the joint entropy $S(AB) = -\text{Tr}(\rho^{AB} \log \rho^{AB})$ by analogy with the corresponding notions for Shannon entropy (cf. Eqs. (19),

⁷Von Neumann first defined this quantity on the basis of a thermodynamic argument in [1955, 379].

(23), (24):

$$(77) \quad S(A|B) = S(A, B) - S(B)$$

$$(78) \quad S(A:B) = S(A) - S(A|B)$$

$$(79) \quad = S(B) - S(B|A)$$

$$(80) \quad = S(A) + S(B) - S(A, B)$$

The joint entropy satisfies the subadditivity inequality:

$$(81) \quad S(A, B) \leq S(A) + S(B)$$

with equality if and only if A and B are uncorrelated, i.e., $\rho^{AB} = \rho^A \otimes \rho^B$.

Now, $S(A|B)$ can be negative, while the conditional Shannon entropy is always positive or zero. Consider, for example, the entangled state $|\Psi\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. Since $|\Psi\rangle$ is a pure state, $S(A, B) = 0$. But $S(A) = S(B) = 1$. So $S(A|B) = S(A, B) - S(A) = -1$. In fact, for a pure state $|\Psi\rangle$ of a composite system AB , $S(A|B) < 0$ if and only if $|\Psi\rangle$ is entangled.

For a composite system AB in a product state $\rho \otimes \sigma$, it follows from the definition of joint entropy that:

$$(82) \quad S(A, B) = S(\rho \otimes \sigma) = S(\rho) + S(\sigma) = S(A) + S(B).$$

If AB is in a pure state $|\Psi\rangle$, it follows from the Schmidt decomposition theorem that $|\Psi\rangle$ can be expressed as

$$(83) \quad |\Psi\rangle = \sum_i \sqrt{p_i} |i\rangle \langle i|$$

from which it follows that

$$(84) \quad \begin{aligned} \rho_A &= \text{Tr}_B(|\psi\rangle\langle\psi|) = \sum_i |i\rangle\langle i| \\ \rho_B &= \text{Tr}_A(|\psi\rangle\langle\psi|) = \sum_i |i\rangle\langle i|; \end{aligned}$$

and so:

$$(85) \quad S(A) = S(B) = - \sum_i p_i \log p_i.$$

Consider a mixed state prepared as a mixture of states ρ_i with weights p_i . It can be shown that

$$(86) \quad S\left(\sum_i p_i \rho_i\right) \leq H(p_i) + \sum_i p_i S(\rho_i)$$

with equality if and only if the states ρ_i have support on orthogonal subspaces (see [Nielsen and Chuang, 2000, Theorem 11.10, p. 518]). The entropy $H(p_i)$ is referred to as the *entropy of preparation* of the mixture ρ .

If the states ρ_i are pure states, then $S(\rho) \leq H(p_i)$. For example, suppose \mathcal{H}^Q is 2-dimensional and $p_1 = p_2 = 1/2$, then $H(p_i) = 1$. So if we had a classical information source producing the symbols 1 and 2 with equal probabilities, no

compression of the information would be possible. However, if the symbols 1 and 2 are encoded as nonorthogonal quantum states $|r_1\rangle$ and $|r_2\rangle$, then $S(\rho) < 1$. As we will see in §3.5, according to Schumacher's source coding theorem, since $S(\rho) < 1$, quantum compression is possible, i.e., we can transmit long sequences of qubits reliably using $S < 1$ qubits per quantum state produced by the source.

Note that if AB is prepared in a mixture of states $\rho_i \otimes |i\rangle\langle i|$ with weights p_i , where the ρ_i are any density operators, not necessarily orthogonal, then it follows from (86), (82), and the fact that $S(|i\rangle\langle i|) = 0$ that

$$\begin{aligned} S\left(\sum_i p_i \rho_i \otimes |i\rangle\langle i|\right) &= H(p_i) + \sum_i p_i S(\rho_i \otimes |i\rangle\langle i|) \\ (87) \qquad \qquad \qquad &= H(p_i) + \sum_i p_i S(\rho_i). \end{aligned}$$

The von Neumann entropy of a mixture of states ρ_i with weights p_i , $\sum p_i \rho_i$, is a concave function of the states in the distribution, i.e.,

$$(88) \quad S\left(\sum_i p_i \rho_i\right) \geq \sum_i p_i S(\rho_i).$$

To see this, consider a composite system AB in the state

$$(89) \quad \rho^{AB} = \sum_i p_i \rho_i \otimes |i\rangle\langle i|.$$

We have

$$(90) \quad S(A) = S\left(\sum_i p_i \rho_i\right)$$

$$(91) \quad S(B) = S\left(\sum_i p_i |i\rangle\langle i|\right) = H(p_i)$$

and

$$(92) \quad S(A, B) = H(p_i) + \sum_i p_i S(\rho_i)$$

by equation (87). By subadditivity $S(A) + S(B) \geq S(A, B)$, so:

$$(93) \quad S\left(\sum_i p_i \rho_i\right) \geq \sum_i p_i S(\rho_i).$$

It turns out that projective measurements always increase entropy, i.e., if $\rho' = \sum_i P_i \rho P_i$, then $S(\rho') \geq S(\rho)$, but generalized measurements can *decrease* entropy. Consider, for example, the generalized measurement on a qubit in the initial state ρ defined by the measurement operators $M_1 = |0\rangle\langle 0|$ and $M_2 = |0\rangle\langle 1|$. (Note that

these operators *do* define a generalized measurement because $M_1^\dagger M_1 + M_2^\dagger M_2 = |0\rangle\langle 0| + |1\rangle\langle 1| = I$.) After the measurement

$$\begin{aligned} \rho' &= |0\rangle\langle 0|\rho|0\rangle\langle 0| + |0\rangle\langle 1|\rho|1\rangle\langle 0| \\ &= \text{Tr}(\rho)|1\rangle\langle 1| \\ (94) \quad &= |1\rangle\langle 1|. \end{aligned}$$

So $S(\rho') = 0 \leq S(\rho)$.

3.3 The ‘No Cloning’ Theorem

In § 3.1 we saw that two nonorthogonal quantum states cannot be reliably distinguished by any measurement. A ‘no cloning’ theorem establishes that nonorthogonal quantum states cannot be copied. To see this, suppose there were a device D that could copy any input quantum state of a system Q with states in \mathcal{H}^Q . Suppose the initial ready state of the device D is $|0\rangle \in \mathcal{H}^D$. Then we require, for any orthonormal set of input states $\{|i\rangle\}$:

$$(95) \quad |i\rangle|0\rangle \xrightarrow{U} |i\rangle|i\rangle$$

where U is the unitary transformation that implements the copying process. By linearity, it then follows that for any input state $\sum_i c_i|i\rangle$:

$$(96) \quad \left(\sum_i c_i|i\rangle\right)|0\rangle \xrightarrow{U} \sum_i c_i|i\rangle|i\rangle$$

But for copying we require that:

$$(97) \quad \left(\sum_i c_i|i\rangle\right)|0\rangle \xrightarrow{U} \left(\sum_i c_i|i\rangle\right)\left(\sum_i c_i|i\rangle\right)$$

and

$$(98) \quad \sum_i c_i|i\rangle|i\rangle \neq \left(\sum_i c_i|i\rangle\right)\left(\sum_i c_i|i\rangle\right) = \sum_{ij} c_i c_j |i\rangle|j\rangle$$

unless $c_i c_j = \delta_{ij}$, which means that the device could not copy any states that are not in the orthonormal set $|i\rangle$.

Alternatively, one might note that if two states $|\psi\rangle$ and $|\phi\rangle$ could be copied, then

$$(99) \quad |\psi\rangle|0\rangle \xrightarrow{U} |\psi\rangle|\psi\rangle$$

$$(100) \quad |\phi\rangle|0\rangle \xrightarrow{U} |\phi\rangle|\phi\rangle$$

Since unitary transformations preserve inner products, we require that

$$(101) \quad \langle\psi|\phi\rangle = \langle\psi|\phi\rangle\langle\psi|\phi\rangle$$

which is possible if and only if $\langle\psi|\phi\rangle = 1$ or 0 . That is: for cloning to be possible, either the states are identical, or they are orthogonal.

The ‘no cloning’ theorem was proved independently by Dieks [1982] and Wootters and Zurek [1982]. An important extension of this result to mixtures is due to Barnum, Caves, Fuchs, Jozsa, and Schumacher [1996a]. In a cloning process, a ready state σ of a system B and the state to be cloned ρ of a system A are transformed into two copies of ρ . In a more general *broadcasting* process, a ready state σ and the state to be broadcast ρ are transformed to a new state ω of AB , where the marginal state ω with respect to both A and B is ρ , i.e.,

$$(102) \rho_A = \text{Tr}_B(\omega) = \rho$$

$$(103) \rho_B = \text{Tr}_A(\omega) = \rho$$

The ‘no cloning’ theorem states that a set of pure states can be cloned if and only if the states are mutually orthogonal. The ‘*no broadcasting*’ theorem states that an arbitrary set of states can be broadcast if and only if they are represented by mutually commuting density operators. Classically, since all pure states are, in a formal sense, orthogonal and all operators (representing real-valued functions on phase space) commute, both cloning and broadcasting are possible. Note that broadcasting reduces to cloning for pure states.

Of course, it is always possible to build a special-purpose device to clone a given (known) quantum state $|\psi\rangle$, because this would simply be a device that prepares the state $|\psi\rangle$. The ‘no cloning’ theorem, from another perspective, is just the statement of the quantum measurement problem (see §7): measurements, in the classical sense of reproducing in a second system a copy of the state of the first system (or, more generally, a ‘pointer state’ that represents the state of the first system), are impossible in quantum mechanics, except for measurements restricted to orthogonal sets of input states.

A modification of the argument leading to Eqs. (99)–(101) shows that no information gain about the identity of nonorthogonal states is possible without disturbing the states. Suppose the device D acts as a measuring device that records some information about the identity of the input state, i.e., the output state of the device is different for different input states $|\psi\rangle, |\phi\rangle$; and that the device does not disturb the input states. Then

$$(104) |\psi\rangle|0\rangle \xrightarrow{U} |\psi\rangle|\psi'\rangle$$

$$(105) |\phi\rangle|0\rangle \xrightarrow{U} |\phi\rangle|\phi'\rangle$$

from which it follows that

$$(106) \langle\psi|\phi\rangle = \langle\psi|\phi\rangle\langle\psi'|\phi'\rangle$$

and so

$$(107) \langle\psi'|\phi'\rangle = 1$$

since $\langle \psi | \phi \rangle \neq 0$ if $|\psi\rangle$ and $|\phi\rangle$ are nonorthogonal. In other words, if there is no disturbance to the nonorthogonal input states, there can be no information gain about the identity of the states. So, for example, an eavesdropper, Eve, could gain no information about the identity of nonorthogonal quantum states communicated between Alice and Bob without disturbing the states, which means that passive eavesdropping is impossible for quantum information.

The observation that a set of pure states can be cloned if and only if they are mutually orthogonal is equivalent to the observation that a set of pure states can be reliably distinguished if and only if they are mutually orthogonal. For if we could distinguish a pair of states $|\psi\rangle$ and $|\phi\rangle$, then we could copy them by simply preparing the states with special-purpose preparation devices for $|\psi\rangle$ and $|\phi\rangle$. And if we could copy the states, then we could prepare as many copies as we liked of each state. Because the product states $|\psi\rangle^{\otimes n}$ and $|\phi\rangle^{\otimes n}$ become orthogonal in the limit as $n \rightarrow \infty$, these states are certainly distinguishable, and so the possibility of cloning the states $|\psi\rangle$ and $|\phi\rangle$ would provide a means of distinguishing them.

Note also that, by a similar argument, cloning would allow different mixtures associated with the same density operator to be distinguished. The equal-weight mixture of qubit states $|\uparrow_z\rangle = |0\rangle, |\downarrow_z\rangle = |1\rangle$ (the eigenstates of the spin observable $Z = \sigma_z$) has the same density operator, $I/2$, as the equal-weight mixture of states $|\uparrow_x\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), |\downarrow_x\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ (the eigenstates of $X = \sigma_x$). Since the cloned states $|\uparrow_x\rangle^{\otimes n}, |\downarrow_x\rangle^{\otimes n}$ become distinguishable from the cloned states $|\uparrow_z\rangle^{\otimes n}, |\downarrow_z\rangle^{\otimes n}$, cloning would allow the two mixtures to be distinguished.

This possibility would also allow superluminal signalling. For suppose Alice and Bob shared the entangled state $\frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)$. If Alice measured X or Z on her qubit, she would steer Bob's qubit into the mixture $\frac{1}{2}|\uparrow_x\rangle\langle\uparrow_x| + \frac{1}{2}|\downarrow_x\rangle\langle\downarrow_x|$ or the mixture $\frac{1}{2}|\uparrow_z\rangle\langle\uparrow_z| + \frac{1}{2}|\downarrow_z\rangle\langle\downarrow_z|$. If Bob could distinguish these mixtures by cloning, in a shorter time than the time taken for light to travel between Alice and Bob, he would be able to ascertain whether Alice measured X or Z , so 1 bit of information would be transferred from Alice to Bob superluminally.

3.4 Accessible Information

The ability to exploit quantum states to perform new sorts of information-processing tasks arises because quantum states have different distinguishability properties than classical states. Of course, it is not the mere lack of distinguishability of quantum states that is relevant here, but the different sort of distinguishability enjoyed by quantum states. This indistinguishability is reflected in the limited *accessibility* of quantum information.

To get a precise handle on this notion of accessibility, consider a classical information source in Shannon's sense, with Shannon entropy $H(X)$. Suppose the source produces symbols represented as the values x (in an alphabet \mathcal{X}) of a random variable X , with probabilities p_x , and that the symbols are encoded as quantum states $\rho_x, x \in X$. The mutual information $H(X : Y)$ (as defined by Eqs. (22), (23), (24)) is a measure of how much information one gains, on average,

about the value of the random variable X on the basis of the outcome Y of a measurement on a given quantum state. The *accessible information* is defined as:

$$(108) \text{Sup } H(X:Y)$$

over all possible measurements.

The *Holevo bound* on mutual information provides an important upper bound to accessible information:

$$(109) H(X:Y) \leq S(\rho) - \sum_x p_x S(\rho_x)$$

where $\rho = \sum_x p_x \rho_x$ and the measurement outcome Y is obtained from a measurement defined by a POVM $\{E_y\}$. Since $S(\rho) - \sum_x p_x S(\rho_x) \leq H(X)$ by Eq. (86), with equality if and only if the states ρ_x have orthogonal support, we have:

$$(110) H(X:Y) \leq H(X)$$

Note that X can be distinguished from Y if and only if $H(X:Y) = H(X)$. If the states ρ_x are orthogonal pure states, then in principle there exists a measurement that will distinguish the states, and for such a measurement $H(X:Y) = H(X)$. In this case, the accessible information is the same as the entropy of preparation of the quantum states, $H(X)$. But if the states are nonorthogonal, then $H(X:Y) < H(X)$ and there is no measurement, even in the generalized sense, that will enable the reliable identification of X .

Note, in particular, that if the values of X are encoded as the pure states of a qubit, then $H(X:Y) \leq S(\rho)$ and $S(\rho) \leq 1$. It follows that at most 1 bit of information can be extracted from a qubit by measurement. If X has k equiprobable values, $H(X) = \log k$. Alice could encode these k values into a qubit by preparing it in an equal-weight mixture of k nonorthogonal pure states, but Bob could only extract at most 1 bit of information about the value of X . For an n -state quantum system associated with an n -dimensional Hilbert space, $S(\rho) \leq \log n$. So even though Alice could encode any amount of information into such an n -state quantum system (by preparing the state as a mixture of nonorthogonal states), the most information that Bob could extract from the state by measurement is $\log n$, which is the same as the maximum amount of information that could be encoded into and extracted from an n -state classical system. It might seem, then, that the inaccessibility of quantum information as quantified by the Holevo bound would thwart any attempt to exploit quantum information to perform nonclassical information-processing tasks. In the following sections, we shall see that this is not the case: surprisingly, the inaccessibility of quantum information can actually be exploited in information-processing tasks that transcend the scope of classical information.

For an insightful derivation of the Holevo bound (essentially reproduced below), see [Nielsen and Chuang, 2000, Theorem 12.1, p. 531]. The basic idea is the following: Suppose Alice encodes the distinguishable symbols of a classical

information source with entropy $H(X)$ as quantum states ρ_x (not necessarily orthogonal). That is, Alice has a quantum system P , the preparation device, with an orthonormal pointer basis $|x\rangle$ corresponding to the values of the random variable X , which are produced by the source with probabilities p_x . The preparation interaction correlates the pointer states $|x\rangle$ with the states ρ_x of a quantum system Q , so that the final state of P and Q after the preparation interaction is:

$$(111) \rho^{PQ} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x.$$

Alice sends the system Q to Bob, who attempts to determine the value of the random variable X by measuring the state of Q . The initial state of P, Q , and Bob's measuring instrument M is:

$$(112) \rho^{PQM} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x \otimes |0\rangle\langle 0|$$

where $|0\rangle\langle 0|$ is the initial ready state of M . Bob's measurement can be described by a quantum operation \mathcal{E} on the Hilbert space $\mathcal{H}^Q \otimes \mathcal{H}^M$ that stores a value of y , associated with a POVM $\{E_y\}$ on \mathcal{H}^Q , in the pointer state $|y\rangle$ of M , i.e., \mathcal{E} is defined for any state $\sigma \in \mathcal{H}^Q$ and initial ready state $|0\rangle \in \mathcal{H}^M$ by:

$$(113) \sigma \otimes |0\rangle\langle 0| \xrightarrow{\mathcal{E}} \sum_y \sqrt{E_y} \sigma \sqrt{E_y} \otimes |y\rangle\langle y|.$$

We have (recall the definition of quantum mutual information in Eqs. (78)–(80)):

$$(114) S(P:Q) = S(P:Q, M)^8$$

because M is initially uncorrelated with PQ and

$$(115) S(P':Q', M') \leq S(P:Q, M)$$

because it can be shown ([Nielsen and Chuang, 2000, Theorem 11.15, p. 522]) that quantum operations never increase mutual information (primes here indicate states after the application of \mathcal{E}). Finally:

$$(116) S(P':Q', M')$$

because discarding systems never increases mutual information ([Nielsen and Chuang, 2000, Theorem 11.15, p. 522]), and so:

$$(117) S(P':M') \leq S(P:Q)$$

which (following some algebraic manipulation) is the statement of the Holevo bound, i.e., (117) reduces to (109).

To see this, note (from (111)) that

$$(118) \rho^{PQ} = \sum_x p_x |x\rangle\langle x| \otimes \rho_x$$

So $S(P) = H(p_x)$, $S(Q) = S(\sum_x p_x \rho_x) = S(\rho)$ and, by (87),

$$(119) \quad S(P, Q) = H(p_x) + \sum_x p_x S(\rho_x)$$

since the states $|x\rangle\langle x| \otimes \rho_x$ have support on orthogonal subspaces in $\mathcal{H}^P \otimes \mathcal{H}^Q$. It follows that

$$(120) \quad \begin{aligned} S(P:Q) &= S(P) + S(Q) - S(P, Q) \\ &= S(\rho) - \sum_x p_x S(\rho_x) \end{aligned}$$

which is the right hand side of the Holevo bound.

For the left hand side:

$$(121) \quad \rho^{P'M'} = \text{Tr}_{Q'}(\rho^{P'Q'M'})$$

$$(122) \quad = \text{Tr}_{Q'}\left(\sum_{xy} p_x |x\rangle\langle x| \otimes \sqrt{E_y} \rho_x \sqrt{E_y} \otimes |y\rangle\langle y|\right)$$

$$(123) \quad = \sum_{xy} p_x \text{Tr}(E_y \rho_x E_y) |x\rangle\langle x| \otimes |y\rangle\langle y|$$

$$(124) \quad = \sum_{xy} p(x, y) |x\rangle\langle x| \otimes |y\rangle\langle y|,$$

since $p(x, y) = p_x p(y | x) = p_x \text{Tr}(\rho_x E_y) = p_x \text{Tr}(\sqrt{E_y} \rho_x \sqrt{E_y})$, and so $S(P' : M') = H(X : Y)$.

The Holevo bound limits the representation of classical bits by qubits. Putting it another way, the Holevo bound characterizes the resource cost of encoding classical bits as qubits: one qubit is necessary and sufficient. Can we represent qubits by bits? If so, what is the cost of a qubit in terms of bits? This question is answered by the following result [Barnum *et al.*, 2001]: A quantum source of nonorthogonal signal states can be compressed with arbitrarily high fidelity to α qubits per signal plus any number of classical bits per signal if and only if α is at least as large as the von Neumann entropy S of the source. This means that a generic quantum source cannot be separated into a classical and quantum part: quantum information cannot be traded for any amount of classical information.

3.5 Quantum Information Compression

As pointed out in §3.2, Shannon's source coding theorem (noiseless channel coding theorem) and the core notion of a typical sequence can be generalized for quantum sources. This was first shown by Jozsa and Schumacher [1994] and Schumacher [1995]. See also [Barnum *et al.*, 1996b].

For a classical information bit source, where the output of the source is given by a random variable X with two possible values x_1, x_2 with probabilities p_1, p_2 , the Shannon entropy of the information produced by the source is $H(X) = H(p_1, p_2)$.

So by Shannon's source coding theorem the information can be compressed and communicated to a receiver with arbitrarily low probability of error by using $H(X)$ bits per signal, which is less than one bit if $p_1 \neq p_2$.

Now suppose the source produces qubit states $|\psi_1\rangle, |\psi_2\rangle$ with probabilities p_1, p_2 . The Shannon entropy of the mixture $\rho = p_1|\psi_1\rangle\langle\psi_1| + p_2|\psi_2\rangle\langle\psi_2|$ is $S(\rho)$. Schumacher's generalization of Shannon's source coding theorem shows that the quantum information encoded in the mixture ρ can be compressed and communicated to a receiver with arbitrarily low probability of error by using $S(\rho)$ qubits per signal, and $S(\rho) < 1$ if the qubit states are nonorthogonal.

Note that the signals considered here are *qubit states*. What Schumacher's theorem shows is that we can reliably communicate the sequence of qubit states produced by the source by sending less than one qubit per signal. Note also that since $S(\rho) < H(p_1, p_2)$ if the qubit states are nonorthogonal, the quantum information represented by the sequence of qubit states can be compressed beyond the classical limit of the classical information associated with the entropy of preparation of ρ (i.e., the Shannon entropy of the random variable whose values are the labels of the qubit states).

Since the individual states in a mixture are not in general distinguishable, there are two distinct sorts of compression applicable to quantum information that do not apply to classical information. In *blind compression*, the sequence of quantum states produced by a source is compressed via a compression scheme that depends only on the identities of the quantum states and their probabilities, i.e., the input to the compression scheme is the density operator associated with the distribution. In *visible compression*, the identity of each individual quantum state produced by the source is assumed to be known, i.e., the input to the compression scheme is an individual quantum state in the sequence produced by the source, and the compression of the state is based on the probability distribution of such states.

An example: an inefficient visible compression scheme of the above qubit source ($|\psi_1\rangle, |\psi_2\rangle$ with probabilities p_1, p_2) would simply involve sending the classical information of the quantum state labels, compressed to $H(p_1, p_2)$ bits per signal, to the receiver, where the original qubit states could then be prepared after decompression of the classical information. This scheme is not optimal by Schumacher's theorem (for nonorthogonal qubit states) because $S(\rho) < H(p_1, p_2)$. Of course, Schumacher's theorem refers to a compression rate of $S(\rho)$ *qubits* per quantum signal, while the application of Shannon's theorem here refers to $H(p_1, p_2)$ *bits* per classical signal. But note that the communication of one classical bit requires the same physical resource as the communication of one qubit, prepared in one of two orthogonal basis states. Note also that sending the (nonorthogonal) qubit states themselves, which would require one qubit per signal, would not convey the identity of the states in the sequence to the transmitter. So the classical information about the individual state labels in the sequence (which would be bounded by $\log n$ per signal if we considered a source producing n qubit states) is really redundant if the aim is to communicate the quantum information associated with the sequence of qubit states.

Remarkably, Schumacher’s theorem shows that the optimal compressibility of the quantum information associated with a sequence of quantum pure states is $S(\rho)$ qubits per signal, for blind *or* visible compression.

To see the general idea, consider a source of (possibly nonorthogonal) qubits $|\psi_1\rangle, |\psi_2\rangle$ with probabilities p_1, p_2 . The density operator of the probability distribution is $\rho = p_1|\psi_1\rangle\langle\psi_1| + p_2|\psi_2\rangle\langle\psi_2|$.

An n -sequence of states produced by the source is represented by a state

$$(125) \quad |\Psi_{i_1\dots i_n}\rangle = |\psi_{i_1}\rangle \dots |\psi_{i_n}\rangle$$

in $\mathcal{H}_2^{\otimes n}$. Each such state has a probability $p_{i_1\dots i_n} = p_{i_1} \dots p_{i_n}$. The n -sequences span the 2^n -dimensional Hilbert space $\mathcal{H}_2^{\otimes n}$, but as $n \rightarrow \infty$ it turns out that the probability of finding an n -sequence in a ‘typical subspace’ (in a measurement, on an n -sequence produced by the source, of the projection operator onto the subspace) tends to 1. That is, for any $\epsilon, \delta > 0$, there is a subspace $\mathcal{T}_\delta^{(n)}$ of dimension between $2^{n(S(\rho)-\delta)}$ and $2^{n(S(\rho)+\delta)}$, with projection operator $P_\delta^{(n)}$, such that:

$$(126) \quad \sum_{\text{all sequences}} p_{i_1\dots i_n} \text{Tr}(|\Psi_{i_1\dots i_n}\rangle\langle\Psi_{i_1\dots i_n}| P_\delta^{(n)}) = \text{Tr}(\rho^{\otimes n} P_\delta^{(n)}) \geq 1 - \epsilon.$$

Here $\rho^{\otimes n} = \rho \otimes \rho \dots \rho$, the n -fold tensor product of ρ with itself, is the density operator of n -sequences of states produced by the source:

$$(127) \quad \rho^{\otimes n} = \sum_{\text{all } n\text{-sequences}} p_{i_1\dots i_n} |\Psi_{i_1\dots i_n}\rangle\langle\Psi_{i_1\dots i_n}|$$

$$(128) \quad = \sum_{\text{all } n\text{-sequences}} p_{i_1} \dots p_{i_n} |\psi_{i_1}\rangle\langle\psi_{i_1}| \otimes \dots \otimes |\psi_{i_n}\rangle\langle\psi_{i_n}|$$

where each state $|\psi_{i_j}\rangle$ is one of k possible states in a d -dimensional Hilbert space. Recall that the statistical properties of such n -sequences of states, for all possible measurements, is given by $\rho^{\otimes n}$ and does not depend on the representation of $\rho^{\otimes n}$ as a particular mixture of states. Since $S(\rho) \leq 1$ for a qubit source, the dimension of $\mathcal{T}_\delta^{(n)}$ decreases exponentially in $\mathcal{H}_2^{\otimes n}$ as $n \rightarrow \infty$, i.e., the typical subspace is exponentially small in $\mathcal{H}_2^{\otimes n}$ for large n .

Note that this does *not* mean that almost all n -sequences of states produced by the source lie in the typical subspace. Rather, almost all n -sequences produced by the source are such that a measurement of $P_\delta^{(n)}$ on the sequence will yield the value 1, i.e., almost all n -sequences produced by the source will answer ‘yes’ in a measurement of the projection operator onto the typical subspace. So, in this sense, most sequences produced by the source will be found to lie in the typical subspace on measurement, and for any subspace \mathcal{V} of dimension less than $2^{n(S(\rho)-\delta)}$ it can be shown that the average probability of finding an n -sequence produced by the source in \mathcal{V} is less than any pre-assigned ϵ for sufficiently large n .

Consider now the general case where the source produces k states $|\psi_1\rangle, \dots, |\psi_k\rangle \in \mathcal{H}_d$ (not necessarily orthogonal) with probabilities p_1, \dots, p_k . Here the density operator associated with the source is $\rho = \sum_{i=1}^k p_i |\psi_i\rangle\langle\psi_i|$. Sequences of length n span a subspace of $d^n = 2^{n \log d}$ dimensions and the typical subspace $\mathcal{T}_\delta^{(n)}$ has dimension between $2^{n(S(\rho)-\delta)}$ and $2^{n(S(\rho)+\delta)}$, which is again exponentially small in $\mathcal{H}_d^{\otimes n}$ because $S(\rho) \leq \log d$.

For comparison with Shannon’s theorem, we write ρ in the spectral representation as:

$$(129) \quad \rho = \sum_x p(x) |x\rangle\langle x|$$

where $\{p(x)\}$ is the set of non-zero eigenvalues of ρ and $\{|x\rangle\}$ is an orthonormal set of eigenstates of ρ . If ρ has eigenvalues $p(x)$ and eigenstates $|x\rangle$, then $\rho^{\otimes n}$ has eigenvalues $p(x_1)p(x_2)\dots p(x_n)$ and eigenstates $|x_1\rangle|x_2\rangle\dots|x_n\rangle$.

A δ -typical state is defined as a state $|x_1\rangle|x_2\rangle\dots|x_n\rangle$ for which the sequence x_1, x_2, \dots, x_n is a δ -typical sequence, in the sense that (cf. Eq. (13)):

$$(130) \quad 2^{-n(S(\rho)+\delta)} < p(x_1 \dots x_n) < 2^{-n(S(\rho)-\delta)}.$$

The δ -typical subspace $\mathcal{T}_\delta^{(n)}$ is the subspace spanned by all the δ -typical states. Denote the projection operator onto $\mathcal{T}_\delta^{(n)}$ by:

$$(131) \quad P_\delta^{(n)} = \sum_{\delta\text{-typical states}} |x_1\rangle\langle x_1| \otimes |x_2\rangle\langle x_2| \dots |x_n\rangle\langle x_n|$$

Then, for a fixed $\delta > 0$, it can be shown that for any $\epsilon > 0$ and sufficiently large n

$$(132) \quad \text{Tr}(P_\delta^{(n)} \rho^{\otimes n}) \geq 1 - \epsilon;$$

and the dimension of $\mathcal{T}_\delta^{(n)}$ ($= \text{Tr}(P_\delta^{(n)})$) satisfies

$$(133) \quad (1 - \epsilon)2^{n(S(\rho)-\delta)} \leq \dim \mathcal{T}_\delta^{(n)} \leq 2^{n(S(\rho)+\delta)}.$$

That is, the dimension of $\mathcal{T}_\delta^{(n)}$ is roughly $2^{nS(\rho)}$, which is exponentially smaller than the dimension of $\mathcal{H}^{\otimes n}$, as $n \rightarrow \infty$.

It follows that the density operator $\rho^{\otimes n}$ can be replaced with a density operator $\tilde{\rho}^{\otimes n}$ with support on the typical subspace (take $\rho^{\otimes n}$ in the spectral representation, where the matrix is diagonal with $2^{n \log d}$ eigenvalues $p(x_1 \dots x_n) = p(x_1) \dots p(x_n)$, and replace all $p(x_1 \dots x_n)$ that do not correspond to typical sequences with zeros).

Before considering a compression/decompression scheme for quantum information, we need a measure of the reliability of such a scheme in terms of the fidelity, as in the case of classical information. The following definition generalizes the classical notion of fidelity in §2 (see [Jozsa, 1998, 70]): If $|\psi\rangle$ is any pure quantum state and ρ any mixed state, the fidelity between ρ and $|\psi\rangle$ is:

$$(134) \quad F(\rho, |\psi\rangle) = \text{Tr}((\rho|\psi\rangle\langle\psi|)) = \langle\psi|\rho|\psi\rangle$$

which is the probability that a measurement of the projection operator $|\psi\rangle\langle\psi|$ in the state ρ yields the outcome 1, i.e., it is the probability that ρ passes a test of being found to be $|\psi\rangle$ on measurement. Note that for a pure state $\rho = |\psi\rangle\langle\psi|$, $F(|\phi\rangle, |\psi\rangle) = |\langle\psi|\phi\rangle|^2$. The fidelity between two mixed states ρ and σ is defined as:⁹

$$(135) \quad F(\rho, \sigma) = \max|\langle\psi|\phi\rangle|^2 = (\text{Tr}(\sqrt{\rho^{1/2}\sigma\rho^{1/2}}))^2.$$

for all purifications $|\psi\rangle$ of ρ and $|\phi\rangle$ of σ . Note that in spite of appearances, $F(\rho, \sigma)$ is symmetric in ρ and σ .

In the case of a source of n -sequences of quantum states $|\Psi_{i_1\dots i_n}\rangle = |\psi_{i_1}\rangle \dots |\psi_{i_n}\rangle$ with prior probabilities $p_{i_1\dots i_n} = p_{i_1} \dots p_{i_n}$, a compression/decompression scheme will in general yield a mixed state $\rho_{i_1\dots i_n}$. The average fidelity of a compression-decompression scheme for an n -sequence of quantum states is defined as:

$$(136) \quad F_n = \sum_{\text{all } n\text{-sequences}} p_{i_1\dots i_n} \text{Tr}(\rho_{i_1\dots i_n} |\Psi_{i_1\dots i_n}\rangle\langle\Psi_{i_1\dots i_n}|)$$

Schumacher's quantum source coding theorem (or quantum noiseless channel coding theorem) for a quantum source that produces quantum states $|\psi_1\rangle \dots |\psi_n\rangle \in \mathcal{H}_d$ with probabilities $p_1 \dots p_n$ (so the density operator corresponding to the output of the source is $\rho = \sum p_i |\psi_i\rangle\langle\psi_i|$), states that

for any $\epsilon, \delta > 0$: (i) there exists a compression/decompression scheme using $S(\rho) + \delta$ qubits per state for n -length sequences produced by the source that can be decompressed by the receiver with fidelity $F_n > 1 - \epsilon$, for sufficiently large n , and (ii) any compression/decompression scheme using $S(\rho) - \delta$ qubits per state for n -length sequences will have a fidelity $F_n < \epsilon$, for sufficiently large n .

A compression/decompression scheme for such a quantum source would go as follows: The transmitter applies a unitary transformation U in $\mathcal{H}_d^{\otimes n}$ (dimension $= d^n = 2^{n \log d}$) which maps any state in the typical subspace onto a linear superposition of sequences of $n \log d$ qubits, where all but the first $nS(\rho)$ qubits are in the state $|0\rangle$, and then transmits the first $nS(\rho)$ qubits to the receiver. So the transmitter compresses $n \log d$ qubits to $nS(\rho)$ qubits. The receiver adds $n \log d - nS(\rho)$ qubits in the state $|0\rangle$ and applies the unitary transformation U^{-1} . Since the initial $nS(\rho)$ qubits will in general be slightly entangled with the remaining $n \log d - nS(\rho)$ qubits, discarding these qubits amounts to tracing over the associated dimensions, so replacing these qubits with the state $|0\rangle$ will produce a mixed state $\tilde{\rho}_n$. The state $U^{-1}\tilde{\rho}_n$ will pass a test of being found to be the original state $|\Psi_{i_1\dots i_n}\rangle$ with fidelity greater than $1 - \epsilon$.

⁹Note that Nielsen and Chuang [2000, 409] define the fidelity $F(\rho, \sigma)$ as the square root of the quantity defined here. If ρ and σ commute, they can be diagonalized in the same basis. The definition then reduces to their definition of the classical fidelity between two probability distributions defined by the eigenvalues of ρ and σ in footnote 3 in §2.1.

4 ENTANGLEMENT ASSISTED QUANTUM COMMUNICATION

In this section I show how entanglement can be exploited as a channel for the reliable transmission of quantum information. I discuss two related forms of entanglement assisted communication: quantum teleportation in §4.1 and quantum dense coding in §4.2.

4.1 *Quantum Teleportation*

As mentioned in §1, Schrödinger introduced the term ‘entanglement’ to describe the peculiar nonlocal correlations of the EPR-state in an extended two-part commentary [1935; 1936] on the Einstein-Podolsky-Rosen argument [Einstein *et al.*, 1935]. Schrödinger regarded entangled states as problematic because they allow the possibility of what he called ‘remote steering’, which he regarded as a mathematical artefact of the Hilbert space theory and discounted as a physical possibility. As it turns out, quantum teleportation is an experimentally confirmed application of remote steering between two separated systems. This was first pointed out in a paper by Bennett, Brassard, Crépeau, Jozsa, Peres, and Wootters [1993] and later experimentally confirmed by several groups using a variety of different techniques [Bouwmeester *et al.*, 1997; Boschi *et al.*, 1998; Furusawa *et al.*, 1998; Nielsen *et al.*, 1998].

In the 1935 paper, Schrödinger considered pure entangled states with a unique biorthogonal decomposition, as well as cases like the EPR-state, where a biorthogonal decomposition is non-unique. He showed that suitable measurements on one system can fix the (pure) state of the entangled distant system, and that this state depends on what observable one chooses to measure, not merely on the outcome of that measurement. In the second paper, he showed that a ‘sophisticated experimenter’, by performing a suitable local measurement on one system, can ‘steer’ the distant system into any mixture of pure states represented by its reduced density operator. So the distant system can be steered (probabilistically, depending on the outcome of the local measurement) into any pure state in the support of the reduced density operator, with a nonzero probability that depends only on the pure state. For a mixture of linearly independent states of the distant system, the steering can be done by performing a local standard projection-valued measurement in a suitable basis. If the states are linearly dependent, the experimenter performs a generalized measurement (associated with a POVM), which amounts to enlarging the experimenter’s Hilbert space by adding an ancilla, so that the dimension of the enlarged Hilbert space is equal to the number of linearly independent states. As indicated in §3.1, Schrödinger’s analysis anticipated the later result by Hughston, Jozsa, and Wootters [1993].

Suppose Alice and Bob, the traditional protagonists in any two-party communication protocol, each holds one of a pair of qubits in the entangled state:

$$(137) |\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B)$$

Bob's qubit separately is in the mixed state $\rho_B = I/2$, which can be interpreted as an equal weight mixture of the orthogonal states $|0\rangle_B, |1\rangle_B$, or, equivalently, as an infinity of other mixtures including, to take a specific example, the equal weight mixture of the four nonorthogonal normalized states:

$$\begin{aligned} |\phi_1\rangle_B &= \alpha|0\rangle_B + \beta|1\rangle_B \\ |\phi_2\rangle_B &= \alpha|0\rangle_B - \beta|1\rangle_B \\ |\phi_3\rangle_B &= \beta|0\rangle_B + \alpha|1\rangle_B \\ |\phi_4\rangle_B &= \beta|0\rangle_B - \alpha|1\rangle_B \end{aligned}$$

That is:

$$(138) \rho_B = I/2 = \frac{1}{4}(|\phi_1\rangle\langle\phi_1| + |\phi_2\rangle\langle\phi_2| + |\phi_3\rangle\langle\phi_3| + |\phi_4\rangle\langle\phi_4|)$$

If Alice measures the observable with eigenstates $|0\rangle_A, |1\rangle_A$ on her qubit A, and Bob measures the corresponding observable on his qubit B, Alice's outcomes will be oppositely correlated with Bob's outcomes (0 with 1, and 1 with 0). If, instead, Alice prepares an ancilla qubit A' in the state $|\phi_1\rangle_{A'} = \alpha|0\rangle_{A'} + \beta|1\rangle_{A'}$ and measures an observable on the pair of qubits $A' + A$ in her possession with eigenstates:

$$(139) |1\rangle = (|0\rangle_{A'}|1\rangle_A - |1\rangle_{A'}|0\rangle_A)/\sqrt{2}$$

$$(140) |2\rangle = (|0\rangle_{A'}|1\rangle_A + |1\rangle_{A'}|0\rangle_A)/\sqrt{2}$$

$$(141) |3\rangle = (|0\rangle_{A'}|0\rangle_A - |1\rangle_{A'}|1\rangle_A)/\sqrt{2}$$

$$(142) |4\rangle = (|0\rangle_{A'}|0\rangle_A + |1\rangle_{A'}|1\rangle_A)/\sqrt{2}$$

(the Bell states defining the Bell basis in $\mathcal{H}^{A'} \otimes \mathcal{H}^A$), she will obtain the outcomes 1, 2, 3, 4 with equal probability of 1/4, and these outcomes will be correlated with Bob's states $|\phi_1\rangle_B, |\phi_2\rangle_B, |\phi_3\rangle_B, |\phi_4\rangle_B$. That is, if Bob checks to see whether his particle is in the state $|\phi_i\rangle_B$ when Alice reports that she obtained the outcome i , he will find that this is always in fact the case. This follows because

$$(143) |\phi_1\rangle_{A'}|\Psi\rangle = \frac{1}{2}(-|1\rangle|\phi_1\rangle_B - |2\rangle|\phi_2\rangle_B + |3\rangle|\phi_3\rangle_B - |4\rangle|\phi_4\rangle_B)$$

In this sense, Alice can steer Bob's particle into any mixture compatible with his density operator $\rho_B = I/2$ by an appropriate local measurement.

What Schrödinger found problematic about entanglement was the possibility of remote steering in the above sense [1935, 556]:

It is rather discomfoting that the theory should allow a system to be steered or piloted into one or the other type of state at the experimenter's mercy in spite of his having no access to it.

Now, remote steering in this probabilistic sense is precisely what makes quantum teleportation possible. Suppose Alice and Bob share a pair of qubits in the entangled state (137) and Alice is given a qubit A' in an *unknown* state $|\phi_1\rangle$ that she

would like to send to Bob. There is no procedure by which Alice can determine the identity of the unknown state, but even if she could, the amount of classical information that Alice would have to send to Bob in order for him to prepare the state $|\phi_1\rangle$ is potentially infinite, since the precise specification of a general normalized qubit state $\alpha|0\rangle + \beta|1\rangle$ requires two real parameters (the number of independent parameters is reduced from four to two because $|\alpha|^2 + |\beta|^2 = 1$ and the overall phase is irrelevant). Alice could send the qubit itself to Bob, but the quantum information in the qubit state might be corrupted by transmission through a possibly noisy environment.

Instead, for the cost of just two bits of classical information, Alice can succeed in communicating the unknown quantum state $|\phi_1\rangle$ to Bob with perfect reliability. What Alice does is to measure the 2-qubit system $A' + A$ in her possession in the Bell basis. Depending on the outcome of her measurement, $i = 1, 2, 3$, or 4 with equal probability, Bob's qubit will be steered into one of the states $|\phi_1\rangle_B$, $|\phi_2\rangle_B$, $|\phi_3\rangle_B$, $|\phi_4\rangle_B$. If Alice communicates the outcome of her measurement to Bob (requiring the transmission of two bits of classical information), Bob can apply one of four local unitary transformations in his Hilbert space to obtain the state $|\phi_1\rangle_B$:

$i = 1$: do nothing, i.e., apply the identity transformation I

$i = 2$: apply the transformation σ_z

$i = 3$: apply the transformation σ_x

$i = 4$: apply the transformation $i\sigma_y$

where $\sigma_x, \sigma_y, \sigma_z$ are the Pauli spin matrices.

The trick that results in the communication of the state $|\phi_1\rangle$ from Alice to Bob, without the qubit A' literally traveling from Alice to Bob, is the ability afforded Alice by the shared entangled state to correlate one of four measurement outcomes (each occurring with probability 1/4) with one of four states that together represent a particular decomposition of Bob's mixed state. The communication of the state of A' is completed by Bob's operation, which requires that Alice sends the two bits of classical information about her measurement outcome to Bob. In the teleportation protocol, the state of the particle A' is destroyed by Alice's measurement and re-created as the state of Bob's particle by Bob's operation — in fact, the systems A and A' end up in an entangled state as the result of Alice's measurement. Note that if the state $|\phi_1\rangle$ of A' were not destroyed there would be two copies of the state, which would violate the quantum 'no cloning' theorem. So neither Alice nor Bob, nor any other party, can gain any information about the identity of the teleported state, because the recording of such information in the state of another quantum system would amount to a partial copying of the information in the teleported state.

Shared entanglement provides a secure and reliable channel for quantum communication. This might be useful for the communication of quantum information

between parties in a cryptographic protocol, or for the transmission of quantum information between the processing components of a quantum computer. It is a feature of an entangled state shared by two parties that the entanglement is not affected by noise in the environment between them. So the reliability of the communication of quantum information by teleportation depends on the reliability of the required classical communication, which can be protected against noise by well-known techniques of error-correcting codes. An entangled state shared by two parties is also unaffected by changes in their relative spatial location. So Alice could teleport a quantum state to Bob without even knowing Bob's location, by broadcasting the two bits of information.

4.2 Quantum Dense Coding

We know from the Holevo bound (see §3.4) that the maximum amount of classical information that can be reliably communicated by encoding the information in the quantum state of a qubit is one bit, even though an arbitrarily large amount of classical information can be encoded in the state of a qubit (by encoding symbols as nonorthogonal quantum states). Quantum dense coding is a procedure, first pointed out by Bennett and Wiesner [1992], for exploiting entanglement to double the amount of classical information that can be communicated by a qubit.

Consider again the Bell states:

$$(144) |1\rangle = (|0\rangle|1\rangle - |1\rangle|0\rangle)/\sqrt{2}$$

$$(145) |2\rangle = (|0\rangle|1\rangle + |1\rangle|0\rangle)/\sqrt{2}$$

$$(146) |3\rangle = (|0\rangle|0\rangle - |1\rangle|1\rangle)/\sqrt{2}$$

$$(147) |4\rangle = (|0\rangle|0\rangle + |1\rangle|1\rangle)/\sqrt{2}$$

Suppose Alice and Bob share a pair of qubits in the state

$$(148) |1\rangle = (|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B)/\sqrt{2}$$

By performing one of four *local* operations on the qubit in her possession defined by the unitary transformations in \mathcal{H}^A :

$$(149) U_1 = I$$

$$(150) U_2 = \sigma_x$$

$$(151) U_3 = \sigma_x$$

$$(152) U_4 = i\sigma_y$$

Alice can transform the state $|1\rangle$ of the qubit pair into any Bell state. For example:

$$(153) I|1\rangle = |1\rangle$$

$$(154) \sigma_z|1\rangle = |2\rangle$$

$$(155) \sigma_x|1\rangle = |3\rangle$$

$$(156) i\sigma_y|1\rangle = |4\rangle$$

So to communicate two classical bits to Bob, Alice applies one of the four operations above to her qubit and sends the qubit to Bob. Bob then performs a measurement on the two qubits in the Bell basis. Since these are orthogonal states, he can distinguish the states and identify Alice's operation.

5 QUANTUM CRYPTOGRAPHY

Over the past few years, quantum cryptography has emerged as perhaps the most successful area of application of quantum information theoretic ideas. The main results have been a variety of provably secure protocols for key distribution, following an original proposal by Bennett and Brassard [1984], and an important 'no go' theorem by Mayers [1996b; 1997] and Lo and Chau [1998]: the impossibility of unconditionally secure two-party quantum bit commitment. The quantum bit commitment theorem generalizes previous results restricted to one-way communication protocols by Mayers [1996a] and by Lo and Chau [1997] and applies to quantum, classical, and quantum-classical hybrid schemes (since classical information, as we have seen, can be regarded as quantum information subject to certain constraints). The restriction to two-party schemes excludes schemes that involve a trusted third-party or trusted channel properties, and the restriction to schemes based solely on the principles of quantum mechanics excludes schemes that exploit special relativistic signalling constraints, or schemes that might involve time machines or the thermodynamics of black holes, etc.

In §5.1, I show how the security of quantum key distribution depends on features of quantum information — no cloning, no information gain without disturbance, entanglement — that prevent an eavesdropper from secretly gaining information about the quantum communication between two parties, i.e., completely undetectable eavesdropping is in principle impossible for quantum communication. In §5.2, I discuss quantum bit commitment and show why unconditionally secure quantum bit commitment is impossible.

5.1 Key Distribution

Quantum Key Distribution Protocols

In a quantum key distribution protocol, the object is for two parties, Alice and Bob, who initially share no information, to exchange information via quantum and classical channels, so as to end up sharing a secret key which they can then use for encryption, in such a way as to ensure that any attempt by an eavesdropper, Eve, to gain information about the secret key will be detected with non-zero probability.

The one-time pad provides a perfectly secure way for Alice and Bob to communicate classical information, but this is also the *only* way that two parties can achieve perfectly security classical communication. The one-time pad is, essentially, a random sequence of bits. If Alice and Bob both have a copy of the one-time pad, Alice can communicate a message to Bob securely by converting

the message to an n -bit binary number (according to some scheme known to both Alice and Bob), and adding (bitwise, modulo 2) the sequence of bits in the binary number to an n -length sequence of bits from the top of the one-time pad. Alice sends the encrypted sequence to Bob, which Bob can then decrypt using the same sequence of bits from his copy of the one-time pad. Since the encrypted message is random, it is impossible for Eve to decrypt the message without a copy of the one-time pad. It is essential to the security of the scheme that the n random bits used to encrypt the message are discarded once the message is transmitted and decrypted, and that a unique random sequence is used for each distinct message — hence the term ‘one-time pad.’

This procedure guarantees perfect privacy, so long as Alice and Bob, and only Alice and Bob, can each be assumed to possess a copy of an arbitrarily long one-time pad. But this means that in order for two parties to communicate secretly, they must already share a secret: the random key. The *key distribution problem* is the problem of how to distribute the key securely in the first place without the key being secretly intercepted during transmission and copied, and the *key storage problem* is the problem of how to store the key securely without it being secretly copied. We would like a procedure that can be guaranteed to be secure against passive eavesdropping, so that Alice and Bob can be confident that their communications are in fact private.

The key idea in *quantum* cryptography is to exploit the indistinguishability of nonorthogonal quantum states, which we saw in §3.3 entails that any information gained by Eve about the identity of such states will introduce some disturbance of the states that can be detected by Alice and Bob, and the ‘no cloning’ theorem, which makes it impossible for Eve to copy quantum communications between Alice and Bob and store them for later analysis (perhaps using, in addition, intercepted classical communications between Alice and Bob).

A large variety of quantum key distribution schemes have been proposed following the original Bennett and Brassard protocol [1984], now known as BB84. The core idea there was for Alice to send Bob a sequence of qubits, prepared with equal probability in one of the states $|0\rangle, |1\rangle, |+\rangle, |-\rangle$, where the pair of orthogonal states $|0\rangle, |1\rangle$ are nonorthogonal to the pair of orthogonal states $|+\rangle, |-\rangle$. Bob measures each qubit randomly in either the basis $|0\rangle, |1\rangle$ or the basis $|+\rangle, |-\rangle$. Following his measurements, he publicly broadcasts the basis he used for each qubit in the sequence, and Alice publicly broadcasts which of these bases is the same as the basis she used to prepare the qubit. Alice and Bob then discard the qubits for which their bases disagree. Since the outcome states of Bob’s measurements are the same as the states Alice prepared, Alice and Bob share a random key on the remaining qubits. They can then sacrifice a portion of these qubits to detect eavesdropping. Alice publicly announces the qubit state she prepared and Bob checks his measurement outcome to confirm this. If they agree on a sufficient number of qubit states (depending on the expected error rate), they conclude that there has been no eavesdropping and use the remaining portion as the secret key. If they don’t agree, they conclude that the qubits have been disturbed by eavesdropping,

in which case they discard all the qubits and begin the procedure again. The actual protocol involves further subtleties in which a perfectly secure secret key is distilled from the ‘raw key’ obtained in this way by techniques of error correction and privacy amplification.

The BB84 scheme solves the key distribution problem, in the sense that Alice and Bob, who initially share no secrets, can end up sharing a secret key via a key distribution protocol that excludes the possibility of eavesdropping, with arbitrarily high reliability (since the length of the sequence of qubits sacrificed to detect eavesdropping can be arbitrarily long). Clearly, it does not solve the key storage problem, since the output of the key distribution protocol is stored as classical information, which is subject to passive eavesdropping.

A scheme proposed by Ekert [1991] allows Alice and Bob to create a shared random key by performing measurements on two entangled qubits. Suppose Alice and Bob share many copies of an entangled pure state of two qubits, say the Bell state $\frac{1}{\sqrt{2}}(|0\rangle|1\rangle - |1\rangle|0\rangle)$ (perhaps emitted by a common source of entangled pairs between Alice and Bob). Alice and Bob agree on three observables that they each measure on their qubits, where the measurements are chosen randomly and independently for each qubit. After a sequence of measurements on an appropriate number of pairs, Alice and Bob announce the directions of their measurements publicly and divide the measurements into two groups: those in which they measured the spin in different directions, and those in which they measured the spin in the same direction. They publicly reveal the outcomes of the first group of measurements and use these to check that the singlet states have not been disturbed by eavesdropping. Essentially, they calculate a correlation coefficient: any attempt by an eavesdropper, Eve, to monitor the particles will disturb the entangled state and result in a correlation coefficient that is bounded by Bell’s inequality and is therefore distinguishable from the correlation coefficient for the entangled state. If Alice and Bob are satisfied that no eavesdropping has occurred, they use the second group of oppositely correlated measurement outcomes as the key.

Quantum Key Distribution via Pre- and Post-Selection

The Ekert scheme solves the key distribution problem as well as the key storage problem, because a new key is generated for each message from the stored entangled states, and there is no information about the key in the entangled states. Here I describe a key distribution protocol that also involves entangled states (see [Bub, 2001b]), but with a different type of test for eavesdropping. Instead of a statistical test based on Bell’s theorem, the test exploits conditional statements about measurement outcomes generated by pre- and post-selected quantum states.

The peculiar features of pre- and post-selected quantum states were first pointed out by Aharonov, Bergmann, and Lebowitz [1964]. If:

- (i) Alice prepares a system in a certain state $|\text{pre}\rangle$ at time t_1 ,
- (ii) Bob measures some observable M on the system at time t_2 ,

- (iii) Alice measures an observable of which $|\text{post}\rangle$ is an eigenstate at time t_3 , and post-selects for $|\text{post}\rangle$,

then Alice can assign probabilities to the outcomes of Bob's M -measurement at t_2 , conditional on the states $|\text{pre}\rangle$ and $|\text{post}\rangle$ at times t_1 and t_3 , respectively, as follows [Aharonov *et al.*, 1964; Vaidman *et al.*, 1987]:

$$(157) \text{prob}(q_k) = \frac{|\langle \text{pre} | P_k | \text{post} \rangle|^2}{\sum_i |\langle \text{pre} | P_i | \text{post} \rangle|^2}$$

where P_i is the projection operator onto the i 'th eigenspace of M . Notice that (157) — referred to as the 'ABL-rule' (Aharonov-Bergmann-Lebowitz rule) in the following — is time-symmetric, in the sense that the states $|\text{pre}\rangle$ and $|\text{post}\rangle$ can be interchanged.

If M is unknown to Alice, she can use the ABL-rule to assign probabilities to the outcomes of various hypothetical M -measurements. The interesting peculiarity of the ABL-rule, by contrast with the usual Born rule for pre-selected states, is that it is possible — for an appropriate choice of observables M, M', \dots , and states $|\text{pre}\rangle$ and $|\text{post}\rangle$ — to assign unit probability to the outcomes of a set of mutually *noncommuting* observables. That is, Alice can be in a position to assert a conjunction of conditional statements of the form: 'If Bob measured M , then the outcome must have been m_i , with certainty, and if Bob measured M' , then the outcome must have been m'_j , with certainty, \dots ', where M, M', \dots are mutually noncommuting observables. Since Bob could only have measured at most one of these noncommuting observables, Alice's conditional information does not, of course, contradict quantum mechanics: she only knows the eigenvalue m_i of an observable M if she knows that Bob in fact measured M .

Vaidman, Aharonov, and Albert [1987] discuss a case of this sort, where the outcome of a measurement of any of the three spin observables $X = \sigma_x, Y = \sigma_y, Z = \sigma_z$ of a spin- $\frac{1}{2}$ particle can be inferred from an appropriate pre- and post-selection. Alice prepares the Bell state

$$(158) |\text{pre}\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle_A |\uparrow_z\rangle_C + |\downarrow_z\rangle_A |\downarrow_z\rangle_C)$$

where $|\uparrow_z\rangle$ and $|\downarrow_z\rangle$ denote the σ_z -eigenstates. Alice sends one of the particles — the channel particle, denoted by the subscript C — to Bob and keeps the ancilla, denoted by A . Bob measures either X, Y , or Z on the channel particle and returns the channel particle to Alice. Alice then measures an observable R on the pair of

particles, where R has the eigenstates (the subscripts A and C are suppressed):

$$(159) |r_1\rangle = \frac{1}{\sqrt{2}}|\uparrow_z\rangle|\uparrow_z\rangle + \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{-i\pi/4})$$

$$(160) |r_2\rangle = \frac{1}{\sqrt{2}}|\uparrow_z\rangle|\uparrow_z\rangle - \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{-i\pi/4})$$

$$(161) |r_3\rangle = \frac{1}{\sqrt{2}}|\downarrow_z\rangle|\downarrow_z\rangle + \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{-i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{i\pi/4})$$

$$(162) |r_4\rangle = \frac{1}{\sqrt{2}}|\downarrow_z\rangle|\downarrow_z\rangle - \frac{1}{2}(|\uparrow_z\rangle|\downarrow_z\rangle e^{-i\pi/4} + |\downarrow_z\rangle|\uparrow_z\rangle e^{i\pi/4})$$

Note that:

$$(163) |\text{pre}\rangle = \frac{1}{\sqrt{2}}(|\uparrow_z\rangle|\uparrow_z\rangle + |\downarrow_z\rangle|\downarrow_z\rangle)$$

$$(164) = \frac{1}{\sqrt{2}}(|\uparrow_x\rangle|\uparrow_x\rangle + |\downarrow_x\rangle|\downarrow_x\rangle)$$

$$(165) = \frac{1}{\sqrt{2}}(|\uparrow_y\rangle|\downarrow_y\rangle + |\downarrow_y\rangle|\uparrow_y\rangle)$$

$$(166) = \frac{1}{2}(|r_1\rangle + |r_2\rangle + |r_3\rangle + |r_4\rangle)$$

Alice can now assign values to the outcomes of Bob's spin measurements via the ABL-rule, whether Bob measured X, Y , or Z , based on the post-selections $|r_1\rangle$, $|r_2\rangle$, $|r_3\rangle$, or $|r_4\rangle$, according to Table 1 (where 0 represents the outcome \uparrow and 1 represents the outcome \downarrow) [Vaidman *et al.*, 1987]:

	σ_x	σ_y	σ_z
r_1	0	0	0
r_2	1	1	0
r_3	0	1	1
r_4	1	0	1

Table 1. $\sigma_x, \sigma_y, \sigma_z$ measurement outcomes correlated with eigenvalues of R

This case can be exploited to enable Alice and Bob to share a private random key in the following way: Alice prepares a certain number of copies (depending on the length of the key and the level of privacy desired) of the Bell state $|\text{pre}\rangle$ in Eq. (158). She sends the channel particles to Bob in sequence and keeps the ancillas. Bob measures X or Z randomly on the channel particles and returns the particles, in sequence, to Alice. Alice then measures the observable R on the ancilla and channel pairs and divides the sequence into two subsequences: the subsequence S_{14} for which she obtained the outcomes r_1 or r_4 , and the subsequence S_{23} for which she obtained the outcomes r_2 or r_3 . The sequence of operations can be implemented on a quantum circuit; see [Metzger, 2000].

To check that the channel particles have not been monitored by Eve, Alice now publicly announces (broadcasts) the indices of the subsequence S_{23} . As is evident from Table 1, for this subsequence she can make conditional statements of the form: ‘For channel particle i , if X was measured, the outcome was 1 (0), and if Z was measured, the outcome was 0 (1)’, depending on whether the outcome of her R -measurement was r_2 or r_3 . She publicly announces these statements as well. If one of these statements, for some index i , does not agree with Bob’s records, Eve must have monitored the i ’th channel particle. (Of course, agreement does not entail that the particle was *not* monitored.)

For suppose Eve measures a different spin component observable than Bob on a channel particle and Alice subsequently obtains one of the eigenvalues r_2 or r_3 when she measures R . Bob’s measurement outcome, either 0 or 1, will be compatible with just one of these eigenvalues, assuming no intervention by Eve. But after Eve’s measurement, both of these eigenvalues will be possible outcomes of Alice’s measurement. So Alice’s retrodictions of Bob’s measurement outcomes for the subsequence S_{23} will not necessarily correspond to Bob’s records. In fact, one can show that if Eve measures X or Z randomly on the channel particles, or if she measures a particular one of the observables X , Y , or Z on the channel particles (the same observable on each particle), the probability of detection in the subsequence S_{23} is $3/8$.

In the subsequence S_{14} , the 0 and 1 outcomes of Bob’s measurements correspond to the outcomes r_1 and r_4 of Alice’s R -measurements. If, following their public communication about the subsequence S_{23} , Alice and Bob agree that there has been no monitoring of the channel particles by Eve, they use the subsequence S_{14} to define a shared raw key.

Note that even a single disagreement between Alice’s retrodictions and Bob’s records is sufficient to reveal that the channel particles have been monitored by Eve. This differs from the eavesdropping test in the Ekert protocol. Note also that Eve only has access to the channel particles, not the particle pairs. So no strategy is possible in which Eve replaces all the channel particles with her own particles and entangles the original channel particles, treated as a single system, with an ancilla by some unitary transformation, and then delays any measurements until after Alice and Bob have communicated publicly. There is no way that Eve can ensure agreement between Alice and Bob without having access to the particle pairs, or without information about Bob’s measurements.

The key distribution protocol as outlined above solves the key distribution problem but not the key storage problem. If Bob actually makes the random choices, measures X or Z , and records definite outcomes for the spin measurements before Alice measures R , as required by the protocol, Bob’s measurement records — stored as classical information — could in principle be copied by Eve without detection. In that case, Eve would know the raw key (which is contained in this information), following the public communication between Alice and Bob to verify the integrity of the quantum communication channel.

To solve the key storage problem, the protocol is modified in the following way:

Instead of actually making the random choice for each channel particle, measuring one of the spin observables, and recording the outcome of the measurement, Bob keeps the random choices and the spin measurements ‘at the quantum level’ until after Alice announces the indices of the subsequence S_{23} of her R measurements. To do this, Bob enlarges the Hilbert space by entangling the quantum state of the channel particle via a unitary transformation with the states of two ancilla particles that he introduces. One particle is associated with a Hilbert space spanned by two eigenstates, $|d_X\rangle$ and $|d_Z\rangle$, of a choice observable or ‘quantum die’ observable D . The other particle is associated with a Hilbert space spanned by two eigenstates, $|p_\uparrow\rangle$ and $|p_\downarrow\rangle$, of a pointer observable P . (See §5.2 for a discussion of how to implement the unitary transformation on the enlarged Hilbert space.)

On the modified protocol (assuming the ability to store entangled states indefinitely), Alice and Bob share a large number of copies of an entangled 4-particle state. When they wish to establish a random key of a certain length, Alice measures R on an appropriate number of particle pairs in her possession and announces the indices of the subsequence S_{23} . Before Alice announces the indices of the subsequence S_{23} , neither Alice nor Bob have stored any classical information. So there is nothing for Eve to copy. After Alice announces the indices of the subsequence S_{23} , Bob measures the observables D and P on his ancillas with these indices and announces the eigenvalue $|p_\uparrow\rangle$ or $|p_\downarrow\rangle$ as the outcome of his X or Z measurement, depending on the eigenvalue of D . If Alice and Bob decide that there has been no eavesdropping by Eve, Bob measures C and P on his ancillas in the subsequence S_{14} . It is easy to see that the ABL-rule applies in this case, just as it applies in the case where Bob actually makes the random choice and actually records definite outcomes of his X or Z measurements before Alice measures R . In fact, if the two cases were not equivalent for Alice — if Alice could tell from her R -measurements whether Bob had actually made the random choice and actually performed the spin measurements, or had merely implemented these actions ‘at the quantum level’ — the difference could be exploited to signal superluminally.

5.2 Bit Commitment

Some History

In a bit commitment protocol, one party, Alice, supplies an encrypted bit to a second party, Bob. The information available in the encrypted bit should be insufficient for Bob to ascertain the value of the bit, but sufficient, together with further information supplied by Alice at a subsequent stage when she is supposed to reveal the value of the bit, for Bob to be convinced that the protocol does not allow Alice to cheat by encrypting the bit in a way that leaves her free to reveal either 0 or 1 at will.

To illustrate the idea, suppose Alice claims the ability to predict the outcomes of elections. To substantiate her claim without revealing valuable information (perhaps to a potential employer, Bob) she suggests the following demonstration: She proposes to record her prediction about whether a certain candidate will win

or lose by writing a 0 (for ‘lose’) or a 1 (for ‘win’) on a note a month before the election. She will then lock the note in a safe and hand the safe to Bob, but keep the key. After the election, she will announce the bit she chose and prove that she in fact made the commitment at the earlier time by handing Bob the key. Bob can then open the safe and read the note.

Obviously, the security of this procedure depends on the strength of the safe walls or the ingenuity of the locksmith. More generally, Alice can send (encrypted) information to Bob that guarantees the truth of an exclusive classical disjunction (equivalent to her commitment to a 0 or a 1) only if the information is biased towards one of the alternative disjuncts (because a classical exclusive disjunction is true if and only if one of the disjuncts is true and the other false). No principle of classical mechanics precludes Bob from extracting this information, so the security of a classical bit commitment protocol can only be a matter of computational complexity.

The question is whether there exists a quantum analogue of this procedure that is *unconditionally secure*: provably secure as a matter of physical law (according to quantum theory) against cheating by either Alice or Bob. Note that Bob can cheat if he can obtain *some* information about Alice’s commitment before she reveals it (which would give him an advantage in repetitions of the protocol with Alice). Alice can cheat if she can delay actually making a commitment until the final stage when she is required to reveal her commitment, or if she can change her commitment at the final stage with a very low probability of detection.

Bennett and Brassard originally proposed a quantum bit commitment protocol in [1984]. The basic idea was to associate the 0 and 1 commitments with two different mixtures represented by the same density operator. As they showed in the same paper, Alice can cheat by adopting an ‘EPR attack’ or cheating strategy: she prepares entangled pairs of qubits, keeps one of each pair (the ancilla) and sends the second qubit (the channel particle) to Bob. In this way she can fake sending one of two equivalent mixtures to Bob and reveal either bit at will at the opening stage by effectively steering Bob’s particle into the desired mixture by an appropriate measurement. Bob cannot detect this cheating strategy.

In a later paper, Brassard, Crépeau, Jozsa, and Langlois [1993] proposed a quantum bit commitment protocol that they claimed to be unconditionally secure. The BCJL scheme was first shown to be insecure by Mayers [1996a]. Subsequently, Mayers [1996b; 1997] and Lo and Chau [1997; 1998] independently showed that the insight of Bennett and Brassard in [1984] can be extended to a proof that a generalized version of the EPR cheating strategy can always be applied, if the Hilbert space is enlarged in a suitable way by introducing additional ancilla particles.

The impossibility of unconditionally secure quantum bit commitment came as something of a surprise to the community of quantum cryptologists and has profound consequences. Indeed, it would not be an exaggeration to say that the significance of the quantum bit commitment theorem for our understanding of quantum mechanics is comparable to Bell’s theorem [Bell, 1964]. Brassard and Fuchs have speculated ([Brassard, 2000; Fuchs, 1997; Fuchs, 2000; Fuchs and Ja-

cobs, 2002]) that quantum mechanics can be derived from two postulates about quantum information: the possibility of secure key distribution and the impossibility of secure bit commitment. We shall see in §7 what this means for the foundations of quantum mechanics.

Perhaps because of the simplicity of the proof and the universality of the claim, the quantum bit commitment theorem is continually challenged in the literature, on the basis that the proof does not cover all possible procedures that might be exploited to implement quantum bit commitment (see, e.g., Yuen [2005]). There seems to be a general feeling that the theorem is ‘too good to be true’ and that there must be a loophole.

In fact, there is no loophole. While Kent [1999a; 1999b] has shown how to implement a secure classical bit commitment protocol by exploiting relativistic signalling constraints in a timed sequence of communications between verifiably separated sites for both Alice and Bob, and Hardy and Kent [2004] and Aharonov, Ta-Shma, Vazirani, and Yao [2005] have investigated the security of ‘cheat-sensitive’ or ‘weak’ versions of quantum bit commitment, these results are not in conflict with the quantum bit commitment theorem. In a bit commitment protocol as usually understood, there is a time interval of arbitrary length, where no information is exchanged, between the end of the commitment stage of the protocol and the opening or unveiling stage, when Alice reveals the value of the bit. Kent’s ingenious scheme effectively involves a third stage between the commitment state and the unveiling stage, in which information is exchanged between Bob’s sites and Alice’s sites at regular intervals until one of Alice’s sites chooses to unveil the originally committed bit. At this moment of unveiling the protocol is not yet complete, because a further sequence of unveilings is required between Alice’s sites and corresponding sites of Bob before Bob has all the information required to verify the commitment at a single site. If a bit commitment protocol is understood to require an arbitrary amount of free time between the end of the commitment stage and the opening stage (in which no step is to be executed in the protocol), then the quantum bit commitment theorem covers protocols that exploit special relativistic signalling constraints.¹⁰

A Key Observation

The crucial insight underlying the proof of the quantum bit commitment theorem is that any step in a quantum bit commitment protocol that requires Alice or Bob to make a definite choice (whether to perform one of a number of alternative measurements, or whether to implement one of a number of alternative unitary transformations) can always be replaced by an EPR cheating strategy in the generalized sense, assuming that Alice and Bob are both equipped with quantum computers. That is, a classical disjunction over definite possibilities — this operation *or* that operation — can always be replaced by a quantum entanglement and a subsequent measurement (perhaps at a more convenient time for the cheater)

¹⁰I am indebted to Dominic Mayers for clarifying this point.

in which one of the possibilities becomes definite. Essentially, the classical disjunction is replaced by a quantum disjunction. This cheating strategy cannot be detected. Similarly, a measurement can be ‘held at the quantum level’ without detection: instead of performing the measurement and obtaining a definite outcome as one of a number of possible outcomes, a suitable unitary transformation can be performed on an enlarged Hilbert space, in which the system is entangled with a ‘pointer’ ancilla in an appropriate way, and the procedure of obtaining a definite outcome can be delayed. The key point is the possibility of keeping the series of transactions between Alice and Bob at the quantum level by enlarging the Hilbert space, until the final exchange of classical information when Alice reveals her commitment.

Any quantum bit commitment scheme will involve a series of transactions between Alice and Bob, where a certain number, n , of quantum systems — the ‘channel particles’ — are passed between them and subjected to various quantum operations (unitary transformations, measurements, etc.), possibly chosen randomly. These operations can always be replaced, without detection, by entangling a channel particle with one or more ancilla particles that function as ‘pointer’ particles for measurements or ‘die’ particles for random choices. In effect, this is the (generalized) EPR cheating strategy.

To illustrate: Suppose, at a certain stage of a quantum bit commitment protocol, that Bob is required to make a random choice between measuring one of two observables, X or Y , on each channel particle he receives from Alice. For simplicity, assume that X and Y each have two eigenvalues, x_1, x_2 and y_1, y_2 . After recording the outcome of the measurement, Bob is required to return the channel particle to Alice. When Alice receives the i ’th channel particle she sends Bob the next channel particle in the sequence. We may suppose that the measurement outcomes that Bob records form part of the information that enables him to confirm Alice’s commitment, once she discloses it (together with further information), so he is not required to report his measurement outcomes to Alice until the final stage of the protocol when she reveals her commitment.

Instead of following the protocol, Bob can construct a device that entangles the input state $|\psi\rangle_C$ of a channel particle with the initial states, $|d_0\rangle_B$ and $|p_0\rangle_B$, of two ancilla particles that he introduces, the first of which functions as a ‘quantum die’ for the random choice and the second as a ‘quantum pointer’ for the measurement. It is assumed that Bob’s ability to construct such a device — in effect, a special purpose quantum computer — is restricted only by the laws of quantum mechanics.

The entanglement is implemented by a unitary transformation in the following way:¹¹ Define two unitary transformations, U_X and U_Y , that implement the X and Y measurements ‘at the quantum level’ on the tensor product of the Hilbert space of the channel particle, \mathcal{H}_C , and the Hilbert space of Bob’s pointer ancilla,

¹¹Note that there is no loss of generality in assuming that the channel particle is in a pure state. If the channel particle is entangled with Alice’s ancillas, the device implements the entanglement via the transformation $I \otimes \dots$, where I is the identity operator in the Hilbert space of Alice’s ancillas.

\mathcal{H}_{B_P} :

$$(167) \quad \begin{aligned} |x_1\rangle_C |p_0\rangle_B &\xrightarrow{U_X} |x_1\rangle_C |p_1\rangle_B \\ |x_2\rangle_C |p_0\rangle_B &\xrightarrow{U_X} |x_2\rangle_C |p_2\rangle_B \end{aligned}$$

and

$$(168) \quad \begin{aligned} |y_1\rangle_C |p_0\rangle_B &\xrightarrow{U_Y} |y_1\rangle_C |p_1\rangle_B \\ |y_2\rangle_C |p_0\rangle_B &\xrightarrow{U_Y} |y_2\rangle_C |p_2\rangle_B \end{aligned}$$

so that

$$(169) \quad |\psi\rangle_C |p_0\rangle_B \xrightarrow{U_X} \langle x_1|\psi\rangle |x_1\rangle_C |p_1\rangle_B + \langle x_2|\psi\rangle |x_2\rangle_C |p_2\rangle_B$$

and

$$(170) \quad |\psi\rangle_C |p_0\rangle_B \xrightarrow{U_Y} \langle y_1|\psi\rangle |y_1\rangle_C |p_1\rangle_B + \langle y_2|\psi\rangle |y_2\rangle_C |p_2\rangle_B$$

The random choice is defined similarly by a unitary transformation V on the tensor product of the Hilbert space of Bob's die ancilla, \mathcal{H}_{B_D} , and the Hilbert space $\mathcal{H}_C \otimes \mathcal{H}_{B_P}$. Suppose $|d_X\rangle$ and $|d_Y\rangle$ are two orthogonal states in \mathcal{H}_{B_D} and that $|d_0\rangle = \frac{1}{\sqrt{2}}|d_X\rangle + \frac{1}{\sqrt{2}}|d_Y\rangle$. Then (suppressing the obvious subscripts) V is defined by:

$$(171) \quad \begin{aligned} |d_X\rangle \otimes |\psi\rangle |p_0\rangle &\xrightarrow{V} |d_X\rangle \otimes U_X |\psi\rangle |p_0\rangle \\ |d_Y\rangle \otimes |\psi\rangle |p_0\rangle &\xrightarrow{V} |d_Y\rangle \otimes U_Y |\psi\rangle |p_0\rangle \end{aligned}$$

so that

$$(172) \quad |d_0\rangle \otimes |\psi\rangle |p_0\rangle \xrightarrow{V} \frac{1}{\sqrt{2}} |d_X\rangle \otimes U_X |\psi\rangle |p_0\rangle + \frac{1}{\sqrt{2}} |d_Y\rangle \otimes U_Y |\psi\rangle |p_0\rangle$$

where the tensor product symbol has been introduced selectively to indicate that U_x and U_y are defined on $\mathcal{H}_C \otimes \mathcal{H}_{B_P}$.

If Bob were to actually choose the observable X or Y randomly, and actually perform the measurement and obtain a particular eigenvalue, Alice's density operator for the channel particle would be:

$$(173) \quad \begin{aligned} &\frac{1}{2} (|\langle x_1|\psi\rangle|^2 |x_1\rangle\langle x_1| + |\langle x_2|\psi\rangle|^2 |x_2\rangle\langle x_2|) \\ &+ \frac{1}{2} (|\langle y_1|\psi\rangle|^2 |y_1\rangle\langle y_1| + |\langle y_2|\psi\rangle|^2 |y_2\rangle\langle y_2|) \end{aligned}$$

assuming that Alice does not know what observable Bob chose to measure, nor what outcome he obtained. But this is precisely the same density operator generated by tracing over Bob's ancilla particles for the state produced in (172). In other words, the density operator for the channel particle is the same for Alice,

whether Bob randomly chooses which observable to measure and actually performs the measurement, or whether he implements an EPR cheating strategy with his two ancillas that produces the transition (172) on the enlarged Hilbert space.

If Bob is required to eventually report what measurement he performed and what outcome he obtained, he can at that stage measure the die ancilla for the eigenstate $|d_X\rangle$ or $|d_Y\rangle$, and then measure the pointer ancilla for the eigenstate $|p_1\rangle$ or $|p_2\rangle$. In effect, if we consider the ensemble of possible outcomes for the two measurements, Bob will have converted the ‘improper’ mixture generated by tracing over his ancillas to a ‘proper’ mixture. But the difference between a proper and improper mixture is undetectable by Alice since she has no access to Bob’s ancillas, and it is only by measuring the composite system consisting of the channel particle together with Bob’s ancillas that Alice could ascertain that the channel particle is entangled with the ancillas.

In fact, if it were possible to distinguish between a proper and improper mixture, it would be possible to signal superluminally: Alice could know instantaneously whether or not Bob performed a measurement on his ancillas by monitoring the channel particles in her possession. Note that it makes no difference whether Bob or Alice measures first, since the measurements are of observables in different Hilbert spaces, which therefore commute.

Clearly, a similar argument applies if Bob is required to choose between alternative unitary operations at some stage of a bit commitment protocol. Perhaps less obviously, an EPR cheating strategy is also possible if Bob is required to perform a measurement or choose between alternative operations on channel particle $i + 1$, conditional on the outcome of a prior measurement on channel particle i , or conditional on a prior choice of some operation from among a set of alternative operations. Of course, if Bob is in possession of all the channel particles at the same time, he can perform an entanglement with ancillas on the entire sequence, considered as a single composite system. But even if Bob only has access to one channel particle at a time (which he is required to return to Alice after performing a measurement or other operation before she sends him the next channel particle), he can always entangle channel particle $i + 1$ with the ancillas he used to entangle channel particle i .

For example, suppose Bob is presented with two channel particles in sequence. He is supposed to decide randomly whether to measure X or Y on the first particle, perform the measurement, and return the particle to Alice. After Alice receives the first particle, she sends Bob the second particle. If Bob measured X on the first particle and obtained the outcome x_1 , he is supposed to measure X on the second particle; if he obtained the outcome x_2 , he is supposed to measure Y on the second particle. If he measured Y on the first particle and obtained the outcome y_1 , he is supposed to apply the unitary transformation U_1 to the second particle; if he obtained the outcome y_2 , he is supposed to apply the unitary transformation U_2 . After performing the required operation, he is supposed to return the second particle to Alice.

It would seem at first sight that Bob has to actually perform a measurement

on the first channel particle and obtain a particular outcome before he can apply the protocol to the second particle, given that he only has access to one channel particle at a time, so an EPR cheating strategy is excluded. But this is not so. Bob's strategy is the following: He applies the EPR strategy discussed above for two alternative measurements to the first channel particle. For the second channel particle, he applies the following unitary transformation on the tensor product of the Hilbert spaces of his ancillas and the channel particle, where the state of the second channel particle is denoted by $|\phi\rangle$, and the state of the pointer ancilla for the second channel particle is denoted by $|q_0\rangle$ (a second die particle is not required):

$$\begin{aligned}
 & |d_X\rangle|p_1\rangle|\phi\rangle|q_0\rangle \xrightarrow{U_C} |d_X\rangle|p_1\rangle \otimes U_X|\phi\rangle|q_0\rangle \\
 & |d_X\rangle|p_2\rangle|\phi\rangle|q_0\rangle \xrightarrow{U_C} |d_X\rangle|p_2\rangle \otimes U_Y|\phi\rangle|q_0\rangle \\
 & |d_Y\rangle|p_1\rangle|\phi\rangle|q_0\rangle \xrightarrow{U_C} |d_Y\rangle|p_1\rangle \otimes U_1|\phi\rangle|q_0\rangle \\
 (174) \quad & |d_Y\rangle|p_2\rangle|\phi\rangle|q_0\rangle \xrightarrow{U_C} |d_Y\rangle|p_2\rangle \otimes U_2|\phi\rangle|q_0\rangle
 \end{aligned}$$

Proof of the Quantum Bit Commitment Theorem

Since an EPR cheating strategy can always be applied without detection, the proof of the quantum bit commitment theorem assumes that at the end of the commitment stage the composite system consisting of Alice's ancillas, the n channel particles, and Bob's ancillas will be represented by some composite entangled state $|0\rangle$ or $|1\rangle$, depending on Alice's commitment,¹² on a Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, where \mathcal{H}_A is the Hilbert space of the particles in Alice's possession at that stage (Alice's ancillas and the channel particles retained by Alice, if any), and \mathcal{H}_B is the Hilbert space of the particles in Bob's possession at that stage (Bob's ancillas and the channel particles retained by Bob, if any).

Now, the density operators $W_B(0)$ and $W_B(1)$, characterizing the information available to Bob for the two alternative commitments, are obtained by tracing the states $|0\rangle$ and $|1\rangle$ over \mathcal{H}_A . If these density operators are the same, then Bob will be unable to distinguish the 0-state from the 1-state without further information from Alice. In this case, the protocol is said to be 'concealing.' What the proof establishes, by an application of the biorthogonal decomposition theorem, is that if $W_B(0) = W_B(1)$ then there exists a unitary transformation in \mathcal{H}_A that will transform $|0\rangle$ to $|1\rangle$. That is, if the protocol is 'concealing' then it cannot be 'binding' on Alice: she can always follow the protocol (with appropriate substitutions of an EPR strategy) to establish the state $|0\rangle$. At the final stage when she is required to reveal her commitment, she can choose to reveal the alternative commitment, depending on circumstances, by applying a suitable unitary transformation in her own Hilbert space to transform $|0\rangle$ to $|1\rangle$ without Bob being able to detect this

¹²More precisely, depending on whether Alice intends to reveal 0 or 1 — since we are assuming that Alice will apply an EPR cheating strategy whenever this is relevant.

move. So either Bob can cheat by obtaining some information about Alice's choice before she reveals her commitment, or Alice can cheat.

The essentials of the proof can be sketched as follows: In the Schmidt decomposition, the states $|0\rangle$ and $|1\rangle$ can be expressed as:

$$\begin{aligned} |0\rangle &= \sum_i \sqrt{p_i} |a_i\rangle |b_i\rangle \\ (175) \quad |1\rangle &= \sum_j \sqrt{p'_j} |a'_j\rangle |b'_j\rangle \end{aligned}$$

where $\{|a_i\rangle\}, \{|a'_j\rangle\}$ are two orthonormal sets of states in \mathcal{H}_A , and $\{|b_i\rangle\}, \{|b'_j\rangle\}$ are two orthonormal sets in \mathcal{H}_B .

The density operators $W_B(0)$ and $W_B(1)$ are defined by:

$$\begin{aligned} W_B(0) &= \text{Tr}_A |0\rangle\langle 0| = \sum_i p_i |b_i\rangle\langle b_i| \\ (176) \quad W_B(1) &= \text{Tr}_A |1\rangle\langle 1| = \sum_j p'_j |b'_j\rangle\langle b'_j| \end{aligned}$$

Bob can't cheat if and only if $W_B(0) = W_B(1)$. Now, by the spectral theorem, the decompositions:

$$\begin{aligned} W_B(0) &= \sum_i p_i |b_i\rangle\langle b_i| \\ W_B(1) &= \sum_j p'_j |b'_j\rangle\langle b'_j| \end{aligned}$$

are unique for the nondegenerate case, where the p_i are all distinct and the p'_j are all distinct. The condition $W_B(0) = W_B(1)$ implies that for all k :

$$\begin{aligned} p_i &= p'_i \\ (177) \quad |b_i\rangle &= |b'_i\rangle \end{aligned}$$

and so

$$\begin{aligned} |0\rangle &= \sum_i \sqrt{p_i} |a_i\rangle |b_i\rangle \\ (178) \quad |1\rangle &= \sum_i \sqrt{p_i} |a'_i\rangle |b_i\rangle \end{aligned}$$

It follows that there exists a unitary transformation $U \in \mathcal{H}_A$ such that

$$(179) \quad \{|a_k\rangle\} \xrightarrow{U} \{|a'_k\rangle\}$$

and hence

$$(180) \quad |0\rangle \xrightarrow{U} |1\rangle$$

As we shall see in §5.2, instead of transforming $|0\rangle$ to $|1\rangle$ by a unitary transformation, Alice could achieve the same effect by preparing the state $|0\rangle$ and measuring in either of two bases, depending on whether she intends to reveal 0 or 1.

The degenerate case can be handled in a similar way. Suppose that $p_1 = p_2 = p'_1 = p'_2 = p$. Then $|b_1\rangle, |b_2\rangle$ and $|b'_1\rangle, |b'_2\rangle$ span the same subspace \mathcal{H} in \mathcal{H}_B , and hence (assuming the coefficients are distinct for $k > 2$):

$$\begin{aligned} |0\rangle &= \sqrt{p}(|a_1\rangle|b_1\rangle + |a_2\rangle|b_2\rangle) + \sum_{k>2} \sqrt{p_k}|a_k\rangle|b_k\rangle \\ |1\rangle &= \sqrt{p}(|a'_1\rangle|b'_1\rangle + |a'_2\rangle|b'_2\rangle) + \sum_{k>2} \sqrt{p_k}|a'_k\rangle|b_k\rangle \\ (181) \quad &= \sqrt{p}(|a''_1\rangle|b_1\rangle + |a''_2\rangle|b_2\rangle) + \sum_{k>2} \sqrt{p_k}|a'_k\rangle|b_k\rangle \end{aligned}$$

where $|a''_1\rangle, |a''_2\rangle$ are orthonormal states spanning \mathcal{H} . Since $\{|a''_1\rangle, |a''_2\rangle, |a_3\rangle, \dots\}$ is an orthonormal set in \mathcal{H}_A , there exists a unitary transformation in \mathcal{H}_A that transforms $\{|a_k\rangle; k = 1, 2, 3, \dots\}$ to $\{|a''_1\rangle, |a''_2\rangle, |a'_3\rangle, \dots\}$, and hence $|0\rangle$ to $|1\rangle$.

The extension of the theorem to the nonideal case, where $W_B(0) \approx W_B(1)$, so that there is a small probability that Bob could distinguish the alternative commitments, shows that Alice has a correspondingly large probability of cheating successfully: there exists a unitary transformation U in \mathcal{H}_A that will transform $W_B(0)$ sufficiently close to $W_B(1)$ so that Alice can reveal whichever commitment she chooses, with a corresponding small probability of Bob being able to detect this move.

How the Theorem Works: An Example

The following example by Asher Peres (private communication) is a beautiful illustration of how the theorem works. (My analysis of the example owes much to correspondence with Adrian Kent and Dominic Mayers.)

Suppose Alice is required to send Bob a channel particle C in an equal weight mixture of the qubit states:

$$(182) \quad |c_0\rangle = |0\rangle$$

$$(183) \quad |c_2\rangle = -\frac{1}{2}|0\rangle + \frac{\sqrt{3}}{2}|1\rangle$$

$$(184) \quad |c_4\rangle = -\frac{1}{2}|0\rangle - \frac{\sqrt{3}}{2}|1\rangle$$

if she commits to 0, and an equal weight mixture of the qubit states:

$$(185) \quad |c_1\rangle = |1\rangle$$

$$(186) \quad |c_3\rangle = \frac{\sqrt{3}}{2}|0\rangle - \frac{1}{2}|1\rangle$$

$$(187) \quad |c_5\rangle = -\frac{\sqrt{3}}{2}|0\rangle - \frac{1}{2}|1\rangle$$

if she commits to 1. Note that these two mixtures have the same density operator:

$$(188) \rho_0 = \rho_1 = I/2$$

Suppose Alice tries to implement an EPR cheating strategy by preparing the entangled state of a system AC :

$$(189) |0\rangle = \frac{1}{\sqrt{3}}(|a_0\rangle|c_0\rangle + |a_2\rangle|c_2\rangle + |a_4\rangle|c_4\rangle)$$

where $\{|a_0\rangle, |a_2\rangle, |a_4\rangle\}$ is an orthonormal basis in the 3-dimensional Hilbert space \mathcal{H}^A of a suitable ancilla system A . If Alice could transform the state $|0\rangle$ to the state:

$$(190) |1\rangle = \frac{1}{\sqrt{3}}(|a_1\rangle|c_1\rangle + |a_3\rangle|c_3\rangle + |a_5\rangle|c_5\rangle)$$

where $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$ is another orthonormal basis in \mathcal{H}^A , by a local unitary transformation in \mathcal{H}^A , she could delay her commitment to the opening stage. If, at that stage, she decides to commit to 0, she measures the observable with eigenstates $\{|a_0\rangle, |a_2\rangle, |a_4\rangle\}$. If she decides to commit to 1, she performs the local unitary transformation taking the state $|0\rangle$ to the state $|1\rangle$ and measures the observable with eigenstates $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$.

Now, $|0\rangle$ can be expressed as:

$$(191) |0\rangle = \frac{1}{\sqrt{3}} \left(|a_0\rangle \frac{|c_3\rangle - |c_5\rangle}{\sqrt{3}} + |a_2\rangle \frac{|c_1\rangle - |c_3\rangle}{\sqrt{3}} + |a_4\rangle \frac{|c_5\rangle - |c_1\rangle}{\sqrt{3}} \right)$$

$$(192) = \frac{1}{\sqrt{3}} \left(\frac{|a_2\rangle - |a_4\rangle}{\sqrt{3}} |c_1\rangle + \frac{|a_0\rangle - |a_2\rangle}{\sqrt{3}} |c_3\rangle + \frac{|a_4\rangle - |a_0\rangle}{\sqrt{3}} |c_5\rangle \right)$$

In this representation of $|0\rangle$, the factor states $\frac{|a_2\rangle - |a_4\rangle}{\sqrt{3}}$, $\frac{|a_0\rangle - |a_2\rangle}{\sqrt{3}}$, $\frac{|a_4\rangle - |a_0\rangle}{\sqrt{3}}$ in \mathcal{H}^A are not orthogonal — in fact, they are coplanar:

$$(193) |a_0\rangle - |a_2\rangle = -(|a_2\rangle - |a_4\rangle) - (|a_4\rangle - |a_0\rangle)$$

So it seems that there cannot be a suitable unitary transformation that will map $|0\rangle$ to $|1\rangle$ and the EPR strategy is blocked!

Of course, this is not the case. To see that there is such a unitary transformation, note that $|0\rangle$ and $|1\rangle$ can be expressed in the Schmidt decomposition as:

$$(194) |0\rangle = \frac{1}{\sqrt{2}} \left(\frac{2|a_0\rangle - |a_2\rangle - |a_4\rangle}{\sqrt{6}} |c_0\rangle + \frac{|a_2\rangle - |a_4\rangle}{\sqrt{2}} |c_1\rangle \right)$$

$$(195) |1\rangle = \frac{1}{\sqrt{2}} \left(\frac{|a_3\rangle - |a_5\rangle}{\sqrt{2}} |c_0\rangle + \frac{-2|a_1\rangle + |a_3\rangle + |a_5\rangle}{\sqrt{6}} |c_1\rangle \right)$$

Clearly, now, there exists a unitary transformation U in \mathcal{H}^A such that:

$$(196) |0\rangle \xrightarrow{U} |1\rangle$$

It follows that:

$$(197) \{|a_0\rangle, |a_2\rangle, |a_4\rangle\} \xrightarrow{U} \{|a'_0\rangle, |a'_2\rangle, |a'_4\rangle\}$$

where $\{|a'_0\rangle, |a'_2\rangle, |a'_4\rangle\}$ is a basis in \mathcal{H}^A , and so

$$(198) |1\rangle = \frac{1}{\sqrt{3}}(|a'_0\rangle|c_0\rangle + |a'_2\rangle|c_2\rangle + |a'_4\rangle|c_4\rangle)$$

$$(199) = \frac{1}{\sqrt{3}}(|a_1\rangle|c_1\rangle + |a_3\rangle|c_3\rangle + |a_5\rangle|c_5\rangle)$$

So Alice could implement the EPR cheating strategy by preparing the state $|1\rangle$ and measuring in the basis $\{|a'_0\rangle, |a'_2\rangle, |a'_4\rangle\}$ for the 0-commitment, or in the basis $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$ for the 1-commitment. Equivalently, of course, she could prepare the state $|0\rangle$ and measure in two different bases, since the unitary transformation that takes $|1\rangle$ to $|0\rangle$ also takes the basis $\{|a_1\rangle, |a_3\rangle, |a_5\rangle\}$ to the basis $\{|a''_1\rangle, |a''_3\rangle, |a''_5\rangle\}$, and so:

$$(200) |0\rangle = \frac{1}{\sqrt{3}}(|a_0\rangle|c_0\rangle + |a_2\rangle|c_2\rangle + |a_4\rangle|c_4\rangle)$$

$$(201) = \frac{1}{\sqrt{3}}(|a''_1\rangle|c_1\rangle + |a''_3\rangle|c_3\rangle + |a''_5\rangle|c_5\rangle)$$

A calculation shows that:

$$(202) |a''_1\rangle = \frac{1}{3} \left(|a_0\rangle + (1 + \sqrt{3})|a_2\rangle + (1 - \sqrt{3})|a_4\rangle \right)$$

$$(203) |a''_3\rangle = \frac{1}{3} \left((1 + \sqrt{3})|a_0\rangle + (1 - \sqrt{3})|a_2\rangle + |a_4\rangle \right)$$

$$(204) |a''_5\rangle = \frac{1}{3} \left((1 - \sqrt{3})|a_0\rangle + |a_2\rangle + (1 + \sqrt{3})|a_4\rangle \right)$$

In effect, if Alice prepares the entangled state $|0\rangle$ and measures the ancilla A in the $\{|a_0\rangle, |a_2\rangle, |a_4\rangle\}$ basis, she steers the channel particle into a mixture of nonorthogonal states $\{|c_0\rangle, |c_2\rangle, |c_4\rangle\}$. If she measures in the $\{|a''_1\rangle, |a''_3\rangle, |a''_5\rangle\}$ basis, she steers the channel particle into a mixture of nonorthogonal states $\{|c_1\rangle, |c_3\rangle, |c_5\rangle\}$.

It follows that Alice can implement the EPR cheating strategy without performing any unitary transformation — she simply entangles the channel particle with a suitable ancilla particle and performs one of two measurements at the opening stage, depending on her commitment. This shows that the unitary transformation required by the theorem is not in fact required. If a cheating strategy is possible in which Alice, at the opening stage, either makes a measurement on an entangled state for the 0-commitment, or transforms this entangled state to a different state by a local unitary transformation in her Hilbert space and then makes a measurement on the transformed state for the 1-commitment, then an equally good cheating strategy is available in which Alice prepares one entangled state for both commitments, and measures in two alternative bases at the opening stage, depending on her commitment.

A Final Worry Laid to Rest

The heart of the mathematical proof is the Schmidt decomposition theorem. But the essential conceptual insight is the possibility of enlarging the Hilbert space and implementing an EPR strategy without detection.

This raises the following question: Suppose Bob cannot cheat because $W_B(0) = W_B(1)$, so by the theorem there exists a unitary transformation U in \mathcal{H}_A that will transform $|0\rangle$ to $|1\rangle$. Could there be a protocol in which Alice also cannot cheat because, although there exists a suitable unitary transformation U , she cannot know what unitary transformation to apply? This is indeed the case, but only if U depends on Bob's operations, which are unknown to Alice. But then Bob would have to actually make a definite choice or obtain a definite outcome in a measurement, and he could always avoid doing so without detection by applying an EPR strategy.

This raises a further question: How do we know that following an EPR strategy is never disadvantageous to the cheater? If so, Bob might choose to avoid an EPR strategy in a certain situation because it would be disadvantageous to him. Could there be a bit commitment protocol where the application of an EPR strategy by Bob at a certain stage of the protocol would give Alice the advantage, rather than Bob, while conforming to the protocol would ensure that neither party could cheat? If there were such a protocol, then Bob would, in effect, be forced to conform to the protocol and avoid the EPR strategy, and unconditionally secure bit commitment would be possible.

In fact, the impossibility of such a protocol follows from the theorem (see [Bub, 2001a]). Suppose there were such a protocol. That is, suppose that if Bob applies an EPR strategy then $W_B(0) = W_B(1)$, so by the theorem there exists a unitary transformation U in Alice's Hilbert space that will transform $|0\rangle$ to $|1\rangle$. Alice must know this U because it is uniquely determined by Bob's deviation from the protocol according to an EPR strategy that keeps all disjunctions at the quantum level as linear superpositions. Suppose also that if, instead, Bob is honest and follows the protocol (so that there is a definite choice for every disjunction over possible operations or possible measurement outcomes), then $W_B(0) = W_B(1)$, but the unitary transformation in Alice's Hilbert space that allows her to transform $|0\rangle$ to $|1\rangle$ depends on Bob's choices or measurement outcomes, which are unknown to Alice.

The point to note is that the information available in Alice's Hilbert space must be the same whether Bob follows the protocol and makes determinate choices and obtains determinate measurement outcomes before Alice applies the unitary transformation U that transforms $|0\rangle$ to $|1\rangle$, or whether he deviates from the protocol via an EPR strategy in which he implements corresponding entanglements with his ancillas to keep choices and measurement outcomes at the quantum level before Alice applies the transformation U , and only makes these choices and measurement outcomes definite at the final stage of the protocol by measuring his ancillas. There can be no difference for Alice because Bob's measurements on his ancillas and

any measurements or operations that Alice might perform take place in different Hilbert spaces, so the operations commute. If Alice's density operator (obtained by tracing over Bob's ancillas), which characterizes the statistics of measurements that Alice can perform in her part of the universe, were different depending on whether or not Bob actually carried out the required measurements, as opposed to keeping the alternatives at the quantum level by implementing corresponding entanglements with ancillas, then it would be possible to use this difference to signal superluminally. Actual measurements by Bob on his ancillas that selected alternatives in the entanglements as determinate would instantaneously alter the information available in Alice's part of the universe.

It follows that in the hypothetical bit commitment protocol we are considering, the unitary transformation U in Alice's Hilbert space that transforms $|0\rangle$ to $|1\rangle$ must be the same transformation in the honest scenario as in the cheating scenario. But we are assuming that the transformation in the honest scenario is unknown to Alice and depends on Bob's measurement outcomes, while the transformation in the cheating scenario is unique and known to Alice. So there can be no such protocol: the deviation from the protocol by an EPR strategy can never place Bob in a worse position than following the protocol honestly.

The argument can be put formally in terms of the theorem as follows: The cheating scenario produces one of two alternative pure states $|0\rangle_c$ or $|1\rangle_c$ in $\mathcal{H}_A \otimes \mathcal{H}_B$ ('c' for 'cheating strategy'). Since the reduced density operators in \mathcal{H}_B :

$$\begin{aligned} W_B^{(c)}(0) &= Tr_A |0\rangle\langle 0|_c \\ (205) \quad W_B^{(c)}(1) &= Tr_A |1\rangle\langle 1|_c \end{aligned}$$

are required by assumption to be the same:

$$(206) \quad W_B^{(c)}(0) = W_B^{(c)}(1)$$

the states $|0\rangle_c$ and $|1\rangle_c$ can be expressed in biorthogonal decomposition as:

$$\begin{aligned} |0\rangle_c &= \sum_i \sqrt{p_i} |a_i\rangle \langle b_i| \\ (207) \quad |1\rangle_c &= \sum_i \sqrt{p_i} |a'_i\rangle \langle b_i| \end{aligned}$$

where the reduced density operators in \mathcal{H}_A :

$$\begin{aligned} W_A^{(c)}(0) = Tr_B |0\rangle\langle 0|_c &= \sum_i p_i |a_i\rangle \langle a_i| \\ (208) \quad W_A^{(c)}(1) = Tr_B |1\rangle\langle 1|_c &= \sum_i p_i |a'_i\rangle \langle a'_i| \end{aligned}$$

are different:

$$(209) \quad W_A^{(c)}(0) \neq W_A^{(c)}(1)$$

It follows that there exists a unitary operator $U_c \in \mathcal{H}_A$ defined by the spectral representations of $W_A^{(c)}(0)$ and $W_A^{(c)}(1)$:

$$(210) \{|a_i\rangle\} \xrightarrow{U_c} \{|a'_i\rangle\}$$

such that:

$$(211) |0\rangle_c \xrightarrow{U_c} |1\rangle_c$$

The honest scenario produces one of two alternative pure states $|0\rangle_h$ and $|1\rangle_h$ in $\mathcal{H}_A \otimes \mathcal{H}_B$ ('h' for 'honest'), where the pair $\{|0\rangle_h, |1\rangle_h\}$ depends on Bob's choices and the outcomes of his measurements.

By assumption, as in the cheating scenario, the reduced density operators $W_B^{(h)}(0)$ and $W_B^{(h)}(1)$ in \mathcal{H}_B are the same:

$$(212) W_B^{(h)}(0) = W_B^{(h)}(1)$$

which entails the existence of a unitary operator $U_h \in \mathcal{H}_A$ such that:

$$(213) |0\rangle_h \xrightarrow{U_h} |1\rangle_h$$

where U_h depends on Bob's choices and measurement outcomes.

Now, the difference between the honest scenario and the cheating scenario is undetectable in \mathcal{H}_A , which means that the reduced density operators in \mathcal{H}_A are the same in the honest scenario as in the cheating scenario:

$$(214) \begin{aligned} W_A^{(h)}(0) &= W_A^{(c)}(0) \\ W_A^{(h)}(1) &= W_A^{(c)}(1) \end{aligned}$$

Since U_h is defined by the spectral representations of $W_A^{(h)}(0)$ and $W_A^{(h)}(1)$, it follows that $U_h = U_c$. But we are assuming that U_h depends on Bob's choices and measurement outcomes, while U_c is uniquely defined by Bob's EPR strategy, in which there are no determinate choices or measurement outcomes. Conclusion: there can be no bit commitment protocol in which neither Alice nor Bob can cheat if Bob honestly follows the protocol, but Alice can cheat if Bob deviates from the protocol via an EPR strategy. If neither Bob nor Alice can cheat in the honest scenario, then Bob and not Alice must be able to cheat in the cheating scenario.

A similar argument rules out a protocol in which neither party can cheat if Bob is honest (as above), but if Bob follows an EPR strategy, then $W_B(0) \approx W_B(1)$, so Bob has some probability of cheating successfully, but Alice has a greater probability of cheating successfully than Bob. Again, the unitary transformation U_c that would allow Alice to cheat with a certain probability of success if Bob followed an EPR strategy would also have to allow Alice to cheat successfully if Bob were honest. But the supposition is that Alice cannot cheat if Bob is honest, because the unitary transformation U_h in that case depends on Bob's choices and measurement outcomes, which are unknown to Alice. It follows that there can be no such protocol.

So there is no loophole in the proof of the theorem. Unconditionally secure quantum bit commitment (in the sense of the theorem) really is impossible.

6 QUANTUM COMPUTATION

6.1 *The Church-Turing Thesis and Computational Complexity*

The classical theory of computation concerns the question of what can be computed, and how efficiently.

Various formal notions of computability by Alonzo Church, Kurt Gödel, and others can all be shown to be equivalent to Alan Turing's notion of computability by a Turing machine (see, e.g., [Lewis and Papadimitriou, 1981]). A Turing machine is an abstract computational device that can be in one of a finite set of possible states. It has a potentially infinite tape of consecutive cells to store information (0, 1, or blank in each cell) and a movable tape head that reads the information in a cell. Depending on the symbol in a cell and the state of the machine, the tape overwrites the symbol, changes the state, and moves one cell to the right or the left until it finally halts at the completion of the computation. A program for a Turing machine T (e.g., a program that executes a particular algorithm for finding the prime factors of an integer) is a finite string of symbols — which can be expressed as a binary number $b(T)$ — indicating, for each state and each symbol, a new state, new symbol, and head displacement. Turing showed that there exists a *universal Turing machine* U that can simulate the program of any Turing machine T with at most a polynomial slow-down, i.e., if we initialize U with $b(T)$ and the input to T , then U performs the same computation as T , where the number of steps taken by U to simulate each step of T is a polynomial function of $b(T)$. The Church-Turing thesis is the proposal to identify the class of computable functions with the class of functions computable by a universal Turing machine. Equivalently, one could formulate the Church-Turing thesis in terms of decision problems, which have yes-or-no answers (e.g., the problem of determining whether a given number is a prime number).

Intuitively, some computations are harder than others, and some algorithms take more time than others. The computational complexity of an algorithm is measured by the number of steps required by a Turing machine to run through the algorithm. A decision problem is said to be in complexity class **P**, hence *easy* or *tractable* if there exists an algorithm for solving the problem in polynomial time, i.e., in a number of steps that is a polynomial function of the size n of the input (the number of bits required to store the input). A problem is said to be *hard* or *intractable* if there does not exist a polynomial-time algorithm for solving the problem. A problem is in complexity class **EXP** if the most efficient algorithm requires a number of steps that is an exponential function of the size n of the input. The number of steps here refers to the worst-case running time, τ , which is of the order $\mathcal{O}(n^k)$ for a polynomial-time algorithm and of the order $\mathcal{O}(2^n)$ for an exponential-time algorithm.

Note that an exponential-time algorithm could be more efficient than a polynomial-time algorithm for some range of input sizes, so the above terminology should be understood with caution. Consider the following example (taken from [Barenco,

1998, 145]): $\tau_P(n) = 10^{-23}n^{1000} + 10^{23}/n \approx \mathcal{O}(n^{1000})$ because, for sufficiently large n , the polynomial term dominates (i.e., $\tau_P(n) < cn^{1000}$ for a fixed factor c), and $\tau_E(n) = 10^{23}n^{1000} + 10^{-23}2^n \approx \mathcal{O}(2^n)$ because, for sufficiently large n , the exponential term dominates (i.e., $\tau_E(n) < c2^n$ for a fixed factor c). But for small enough values of n , $\tau_E(n) < \tau_P(n)$.

A Turing machine as defined above is a deterministic machine. A *nondeterministic* or *probabilistic* Turing machine makes a random choice between multiple transitions (to a new symbol, new state, and head displacement) for each symbol and each state. For each sequence of choices, the sequence of transitions corresponds to a sequence of steps executed by a deterministic Turing machine. If any of these machines halts, the computation is regarded as completed. Evidently, a nondeterministic Turing machine cannot compute a function that is not computable by a deterministic Turing machine, but it is believed (but not proved) that certain problems can be solved more efficiently by nondeterministic Turing machines than by any deterministic Turing machine. The complexity class **NP** is the class of problems that can be solved in polynomial time by a nondeterministic Turing machine. This is equivalent to the class of problems for which proposed solutions can be verified in polynomial time by a deterministic Turing machine. For example, it is believed (but not proved) that the problem of factoring an integer into its prime factors is a ‘hard’ problem: there is no known polynomial-time algorithm for this problem. However, the problem of checking whether a candidate factor of an integer is indeed a factor can be solved in polynomial time, so factorizability is an **NP** problem.

Clearly $\mathbf{P} \subseteq \mathbf{NP}$, but it is an open problem in complexity theory whether $\mathbf{P} = \mathbf{NP}$. An **NP** problem is said to be **NP**-complete if every **NP** problem can be shown to have a solution with a number of steps that is a polynomial function of the number of steps required to solve the **NP**-complete problem. So if an **NP**-complete problem can be solved in polynomial time, then all **NP** problems can be solved in polynomial time, and $\mathbf{P} = \mathbf{NP}$. The problem of determining whether a Boolean function $f\{0,1\}^n \rightarrow \{0,1\}$ is satisfiable (i.e., whether there is a set of input values for which the function takes the value 1, or equivalently whether there is an assignment of truth values to the atomic sentences of a compound sentence of Boolean logic under which the compound sentence comes out true) is an **NP**-complete problem. Factorizability is an **NP** problem that is not known to be **NP**-complete.

Since a Turing machine can simulate any classical computing device with at most a polynomial slow-down, the complexity classes are the same for any model of computation. For example, a circuit computer computes the value of a function by transforming data stored in an input register, representing the input to the function, via Boolean circuits constructed of elementary Boolean gates connected by wires, to data in an output register representing the value of the function computed. The elementary Boolean gates are 1-bit gates (such as the NOT gate, which transforms 0 to 1, and conversely) and 2-bit gates (such as the AND gate, which takes two input bits to 1 if and only if they are both 1, otherwise to 0),

and it can be shown that a combination of such gates forms a ‘universal set’ that suffices for any transformation of n bits. In fact, it turns out that one of the sixteen possible 2-bit Boolean gates, the NAND gate (or NOT AND gate), which takes two input bits to 0 if and only if they are both 1, forms a universal set by itself.

In a circuit model of a quantum computer, the registers store qubits, which are then manipulated by elementary unitary gates. It can be shown (see [Nielsen and Chuang, 2000, 188]) that a set of single-qubit and two-qubit unitary gates — the CNOT gate, the Hadamard gate, the phase gate, and the $\pi/8$ gate — forms a universal set, in the sense that any unitary transformation of n qubits can be approximated to arbitrary accuracy by a quantum circuit consisting of these gates connected in some combination. The CNOT gate (‘C’ for ‘controlled’) has two input qubits, a ‘control’ qubit and a ‘target’ qubit. The gate functions so as to flip the target qubit if and only if the control qubit is $|1\rangle$. The remaining three gates are single-qubit gates. The Hadamard gate transforms $|0\rangle$ to $(|0\rangle + |1\rangle)/\sqrt{2}$ and $|1\rangle$ to $(|0\rangle - |1\rangle)/\sqrt{2}$ and is sometimes referred to as the ‘square root of NOT’ gate because two successive applications transforms $|0\rangle$ to $|1\rangle$, and conversely. The phase gate leaves $|0\rangle$ unchanged and transforms $|1\rangle$ to $i|1\rangle$. The $\pi/8$ gate leaves $|0\rangle$ unchanged and transforms $|1\rangle$ to $e^{i\pi/4}|1\rangle$. (See [Nielsen and Chuang, 2000, 174] for a discussion and why the $\pi/8$ gate is so named.)

There are other models of quantum computation. In the ‘cluster state’ or ‘one-way’ quantum computer of Raussendorf and Briegel [2001b; 2001a], a fixed multi-qubit entangled state (called a ‘cluster state’), independent of the computation, is prepared. Then a sequence of single-qubit measurements is performed on this state, where the choice of what observables to measure depends on the outcomes of the previous measurements. No unitary transformations are involved. Remarkably, it can be shown that any quantum circuit of unitary gates and measurements can be simulated by a cluster state computer with similar resources of qubits and time (see [Jozsa, 2005; Nielsen, 2003; Nielsen, 2005]).

The interesting question is whether a quantum computer can perform computational tasks that are not possible for a Turing machine, or perform such tasks more efficiently than any Turing machine. Since a Turing machine is defined by its program, and a program can be specified by a finite set of symbols, there are only countably many Turing machines. There are uncountably many functions on the natural numbers, so there are uncountably many uncomputable functions, i.e., functions that are not computable by any Turing machine. A quantum computer cannot compute a function that is not Turing-computable, because a Turing machine can simulate (albeit inefficiently, with an exponential slow-down [Feynman, 1982]) the dynamical evolution of any system, classical or quantum, with arbitrary accuracy. But there are computational tasks that a quantum computer can perform by exploiting entanglement that are impossible for any Turing machine. Recall the discussion of Bell’s counterargument to the EPR argument in §3.1: a quantum computer, but no classical computer, can perform the task of rapidly producing pairs of values (0 or 1) for pairs of input angles at different locations,

with correlations that violate Bell's inequality, where the response time is less than the time taken by light to travel between the locations.

The current interest in quantum computers concerns the question of whether a quantum computer can compute certain Turing-computable functions more *efficiently* than any Turing machine. In the following section, I discuss quantum algorithms that achieve an exponential speed-up over any classical algorithm, or an exponential speed-up over any *known* classical algorithm. The most spectacular of these is Shor's factorization algorithm, and a related algorithm for solving the discrete log problem.¹³

The factorization algorithm has an important practical application to cryptography. Public-key distribution protocols such as RSA [Rivest *et al.*, 1978] (widely used in commercial transactions over the internet, transactions between banks and financial institutions, etc.) rely on factoring being a 'hard' problem. (Preskill [2005] notes that currently the 65-digit factors of a 130-digit integer can be found in about a month using a network of hundreds of work stations implementing the best known classical factoring algorithm (the 'number sieve algorithm'). He estimates that factoring a 400-digit integer would take about 10^{10} years, which is the age of the universe.) To see the idea behind the RSA protocol, suppose Alice wishes to send a secret message to Bob. Bob's public key consists of two large integers, s and c . Alice encrypts the message m (in the form of a binary number) as $e = m^s \bmod c$ and sends the encrypted message to Bob. Bob decrypts the message as $e^t \bmod c$ where t is an integer known only to Bob. The integer t for which $m = e^t \bmod c$ can easily be determined from s and the factors of c , but since $c = pq$ is the product of two large prime numbers known only to Bob, an eavesdropper, Eve, can read the message only if she can factor c into its prime factors. The cleverness of the scheme resides in the fact that no secret key needs to be distributed between Alice and Bob: Bob's key $\{s, c\}$ is public and allows anyone to send encrypted messages to Bob. If a quantum computer could be constructed that implemented Shor's algorithm, key distribution protocols that rely on the difficulty of factoring very large numbers would be insecure.

6.2 Quantum Algorithms

In the following three sections, I look at the information-processing involved in Deutsch's XOR algorithm [1985], Simon's period-finding algorithm [1994; 1997], and Shor's factorization algorithm [1994; 1997] in terms of the difference between the Boolean logic underlying a classical computation and the non-Boolean logic represented by the projective geometry of Hilbert space, in which the subspace structure of Hilbert space replaces the set-theoretic structure of classical logic. The three algorithms all turn out to involve a similar geometric formulation.

Basically, all three algorithms involve the determination of a global property of a function, i.e., a disjunctive property. The disjunction is represented as a sub-

¹³The discrete log of x with respect to a given prime integer p and an integer q coprime to p is the integer r such that $q^r = x \bmod p$. See [Nielsen and Chuang, 2000, 238] for a discussion.

space in an appropriate Hilbert space, and alternative possible disjunctions turn out to be represented as orthogonal subspaces, except for intersections or overlaps. The true disjunction is determined as the subspace containing the state vector via a measurement. The algorithm generally has to be run several times because the state might be found in the overlap region. The essential feature of these quantum computations is that the true disjunction is distinguished from alternative disjunctions without determining the truth values of the disjuncts. In a classical computation, distinguishing the true disjunction would be impossible without the prior determination of the truth values of the disjuncts. More generally, a quantum computer computes a global property of a function without computing information that is redundant quantum mechanically, but essential for a classical computation of the global property.

There are other quantum algorithms besides these three, e.g., Grover's sorting algorithm [1997] which achieves a quadratic speed-up over any classical algorithm. For a discussion, see [Nielsen and Chuang, 2000], [Jozsa, 1999].

Deutsch's XOR Algorithm and the Deutsch-Jozsa Algorithm

Let $B = \{0, 1\}$ be a Boolean algebra (or the additive group of integers mod 2). In Deutsch's XOR problem [1985], we are given a 'black box' or oracle that computes a function $f : B \rightarrow B$ and we are required to determine whether the function is 'constant' (takes the same value for both inputs) or 'balanced' (takes a different value for each input). Classically, the only way to do this would be to consult the oracle twice, for the input values 0 and 1, and compare the outputs.

In a quantum computation of the Boolean function, a unitary transformation $U_f : |x\rangle|y\rangle \rightarrow |x\rangle|y \oplus f(x)\rangle$ corresponding to the 'black box' correlates input values with corresponding output values.¹⁴ The computation proceeds as follows: The input and output registers are 1-qubit registers initialized to the state $|0\rangle|0\rangle$ in a standard basis. A Hadamard transformation is applied to the input register, yielding a linear superposition of states corresponding to the two possible input values 0 and 1, and the transformation U_f is then applied to both registers, yielding the transitions:

$$(215) \quad |0\rangle|0\rangle \xrightarrow{H} \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|0\rangle$$

$$(216) \quad \xrightarrow{U_f} \frac{1}{\sqrt{2}}(|0\rangle|f(0)\rangle + |1\rangle|f(1)\rangle)$$

If the function is constant, the final composite state of both registers is one of

¹⁴Note that two quantum registers are required to compute functions that are not 1-1 by a unitary transformation. Different input values, x and y , to a function f are represented by orthogonal states $|x\rangle, |y\rangle$. So if $f(x) = f(y)$ for some $x \neq y$, the transformation $W_f : |x\rangle \rightarrow |f(x)\rangle$ could not be unitary, because the orthogonal states $|x\rangle, |y\rangle$ would have to be mapped onto the same state by W_f . The ability of unitary transformations, which are reversible, to compute irreversible functions is achieved by keeping a record of the input for each output value of the function.

the two orthogonal states:

$$(217) |c_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|0\rangle)$$

$$(218) |c_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|1\rangle)$$

If the function is balanced, the final composite state is one of the two orthogonal states:

$$(219) |b_1\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle)$$

$$(220) |b_2\rangle = \frac{1}{\sqrt{2}}(|0\rangle|1\rangle + |1\rangle|0\rangle)$$

The states $|c_1\rangle, |c_2\rangle$ and $|b_1\rangle, |b_2\rangle$ span two planes P_c, P_b in $\mathcal{H}^2 \otimes \mathcal{H}^2$, represented by the projection operators:

$$(221) P_c = P_{|c_1\rangle} + P_{|c_2\rangle}$$

$$(222) P_b = P_{|b_1\rangle} + P_{|b_2\rangle}$$

These planes are orthogonal, except for an intersection, so their projection operators commute. The intersection is the line (ray) spanned by the vector¹⁵:

$$(223) \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle) = \frac{1}{\sqrt{2}}(|c_1\rangle + |c_2\rangle) = \frac{1}{\sqrt{2}}(|b_1\rangle + |b_2\rangle)$$

In the ‘prime’ basis spanned by the states $|0'\rangle = H|0\rangle, |1'\rangle = H|1\rangle$ the intersection is the state $|0'\rangle|0'\rangle$, the ‘constant’ plane is spanned by $|0'\rangle|0'\rangle, |0'\rangle|1'\rangle$, and the ‘balanced’ plane is spanned by $|0'\rangle|0'\rangle, |1'\rangle|1'\rangle$. Note that:

$$(224) |0'\rangle|1'\rangle = \frac{1}{\sqrt{2}}(|c_1\rangle - |c_2\rangle)$$

$$(225) |1'\rangle|1'\rangle = \frac{1}{\sqrt{2}}(|b_1\rangle - |b_2\rangle)$$

In the usual formulation of the algorithm, to decide whether the function f is constant or balanced we measure the output register in the prime basis. If the outcome is $0'$ (which is obtained with probability $1/2$, whether the state ends up in the constant plane or the balanced plane), the computation is inconclusive, yielding no information about the function f . If the outcome is $1'$, then we measure the input register. If the outcome of the measurement on the input register is $0'$, the function is constant; if it is $1'$, the function is balanced.

Alternatively — and this will be relevant for the comparison with Simon’s algorithm and Shor’s algorithm — we could measure the observable with eigenstates $|0'0'\rangle, |0'1'\rangle, |1'0'\rangle, |1'1'\rangle$. The final state is in the 3-dimensional subspace orthogonal to the vector $|1'0'\rangle$, either in the constant plane or the balanced plane.

¹⁵Here $|00\rangle = |0\rangle|0\rangle$, etc.

If the state is in the constant plane, we will either obtain the outcome $0'0'$ with probability $1/2$ (since the final state is at an angle $\pi/4$ to $|0'0'\rangle$), in which case the computation is inconclusive, or the outcome $0'1'$ with probability $1/2$. If the state is in the balanced plane, we will again obtain the outcome $0'0'$ with probability $1/2$, in which case the computation is inconclusive, or the outcome $1'1'$ with probability $1/2$. So in either case, with probability $1/2$, we can distinguish in one run of the algorithm between the two quantum disjunctions ‘constant’ and ‘balanced’ represented by the planes:

$$(226) P_c = P_{|0'0'\rangle} \vee P_{|0'1'\rangle}$$

$$(227) P_b = P_{|0'0'\rangle} \vee P_{|1'1'\rangle}$$

without finding out the truth values of the disjuncts in the computation (i.e., whether in the ‘constant’ case the function maps 0 to 0 and 1 to 0 or whether the function maps 0 to 1 and 1 to 1, and similarly in the ‘balanced’ case). Note that we could also apply a Hadamard transformation to the final states of both registers and measure in the computational basis, since $|0'0'\rangle \xrightarrow{H} |00\rangle$, etc.

Deutsch’s XOR algorithm was the first quantum algorithm with a demonstrated speed-up over any classical algorithm performing the same computational task. However, the algorithm has an even probability of failing, so the improvement in efficiency over a classical computation is only achieved if the algorithm succeeds, and even then is rather modest: one run of the quantum algorithm versus two runs of a classical algorithm. The following variation of Deutsch’s algorithm avoids this feature [Cleve *et al.*, 1998].

We begin by initializing the two registers to $|0\rangle$ and $|1\rangle$, respectively (instead of to $|0\rangle$ and $|0\rangle$) and apply a Hadamard transformation to both registers, which yields the transition:

$$(228) |0\rangle|1\rangle \xrightarrow{H} \frac{|0\rangle + |1\rangle}{\sqrt{2}} \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

Since

$$(229) U_f |x\rangle|y\rangle = |x\rangle|y \oplus f(x)\rangle$$

it follows that

$$(230) U_f |x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \begin{cases} |x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} & \text{if } f(x) = 0 \\ -|x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} & \text{if } f(x) = 1 \end{cases}$$

which can be expressed as

$$(231) U_f |x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} = (-1)^{f(x)} |x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

Notice that the value of the function now appears as a phase of the final state of the input register, a feature referred to as ‘phase kickback.’ For the input state $1/\sqrt{2}(|0\rangle + |1\rangle)$, we have:

$$(232) \quad U_f \frac{|0\rangle + |1\rangle}{\sqrt{2}} \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \frac{(-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle}{\sqrt{2}} \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

which can be expressed as:

$$(233) \quad U_f \frac{|0\rangle + |1\rangle}{\sqrt{2}} \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \begin{cases} \pm \frac{|0\rangle + |1\rangle}{\sqrt{2}} \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \pm |0'\rangle |1'\rangle & \text{if } f(0) = f(1) \\ \pm \frac{|0\rangle - |1\rangle}{\sqrt{2}} \frac{|0\rangle - |1\rangle}{\sqrt{2}} = \pm |1'\rangle |1'\rangle & \text{if } f(0) \neq f(1) \end{cases}$$

Instead of the final state of the two registers ending up as one of two orthogonal states in the constant plane, or as one of two orthogonal states in the balanced plane, the final state now ends up as $\pm|0'1'\rangle$ in the constant plane, or as $\pm|1'1'\rangle$ in the balanced plane, and these states can be distinguished because they are orthogonal. So we can decide with certainty whether the function is constant or balanced after only one run of the algorithm. In fact, we can distinguish these two possibilities by simply measuring the input register in the prime basis. Note that if we perform a final Hadamard transformation on the input register (which takes $|0'\rangle$ to $|0\rangle$ and $|1'\rangle$ to $|1\rangle$), we can distinguish the two possibilities by measuring the input register in the computational basis. Note also that the state of the output register is unchanged: at the end of the process it is in the state $|1'\rangle = H|1\rangle$ (as in (228)) and is not measured.

Deutsch’s XOR problem can be generalized to the problem (‘Deutsch’s problem’) of determining whether a Boolean function $f : B^n \rightarrow B$ is constant or whether it is balanced, where it is promised that the function is either constant or balanced. ‘Balanced’ here means that the function takes the values 0 and 1 an equal number of times, i.e., 2^{n-1} times each. The Deutsch-Jozsa algorithm [1992] decides whether f is constant or balanced in one run.

We begin by setting the input n -qubit register to the state $|0\rangle$ (an abbreviation for the state $|0 \cdots 0\rangle = |0\rangle \cdots |0\rangle$) and the output 1-qubit register to the state $|1\rangle$. We apply an n -fold Hadamard transformation to the input register and a Hadamard transformation to the output register, followed by the unitary transformation U_f to both registers, and finally an n -fold Hadamard transformation to the input register.

First note that

$$(234) \quad H|x\rangle = \frac{1}{\sqrt{2}} \sum_{y \in \{0,1\}} (-1)^{xy} |y\rangle$$

so

$$(235) \quad H^{\otimes n} |x_1, \dots, x_n\rangle = \frac{1}{\sqrt{2^n}} \sum_{y_1, \dots, y_n} (-1)^{x_1 y_1 + \dots + x_n y_n} |y_1, \dots, y_n\rangle$$

This can be expressed as:

$$(236) H^{\otimes n}|x\rangle = \frac{1}{\sqrt{2^n}} \sum_{y \in \{0,1\}^n} (-1)^{x \cdot y} |y\rangle$$

where $x \cdot y$ is the bitwise inner product of x and y , mod 2.

The unitary transformations (Hadamard transformation, U_f) yield:

$$(237) |0\rangle^{\otimes n} |1\rangle \xrightarrow{H} \sum_{x \in \{0,1\}^n} \frac{|x\rangle}{\sqrt{2^n}} \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

$$(238) \xrightarrow{U_f} \sum_x \frac{(-1)^{f(x)}}{\sqrt{2^n}} |x\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

$$(239) \xrightarrow{H} \sum_y \sum_x \frac{(-1)^{x \cdot y + f(x)}}{\sqrt{2^n}} \frac{|y\rangle}{\sqrt{2^n}} \frac{|0\rangle - |1\rangle}{\sqrt{2}}$$

Now consider the state of the input register:

$$(240) \sum_y \sum_x \frac{(-1)^{x \cdot y + f(x)}}{2^n} |y\rangle = \sum_x \frac{(-1)^{f(x)}}{2^n} |0 \dots 0\rangle + \dots$$

Note that the coefficient (amplitude) of the state $|0 \dots 0\rangle$ in the linear superposition (238) is $\sum_x \frac{(-1)^{f(x)}}{2^n}$. If f is constant, this coefficient is ± 1 , so the coefficients of the other terms must all be 0. If f is balanced, $f(x) = 0$ for half the values of x and $f(x) = 1$ for the other half, so the positive and negative contributions to the coefficient of $|0 \dots 0\rangle$ cancel to 0. In other words, if f is constant, the state of the input register is $\pm|0 \dots 0\rangle$; if f is balanced, the state is in the orthogonal subspace.

This is the usual way of describing how the algorithm works, which rather obscures the geometric picture. Consider, for simplicity, the case $n = 2$. After the transformation U_f , but before the final Hadamard transformation, the state of the input register is:

$$(241) \pm \frac{1}{2} (|00\rangle + |01\rangle + |10\rangle + |11\rangle)$$

if the function is constant, or a state of the form:

$$(242) \frac{1}{2} (\pm|00\rangle \pm |01\rangle \pm |10\rangle \pm |11\rangle)$$

if the function is balanced, where two of the coefficients are $+1$ and two of the coefficients are -1 . Evidently, there are three (distinct) mutually orthogonal such balanced states, and they are all orthogonal to the constant state. So the three balanced states lie in a 3-dimensional subspace orthogonal to the constant state and can therefore be distinguished from the constant state. The final Hadamard transformation transforms the constant state:

$$(243) \pm \frac{1}{2} (|00\rangle + |01\rangle + |10\rangle + |11\rangle) \xrightarrow{H} \pm |00\rangle$$

and the three balanced states to states in the 3-dimensional subspace orthogonal to $|00\rangle$. So to decide whether the function is constant or balanced we need only measure the input register and check whether it is in the state $|00\rangle$.

Simon's Algorithm

The problem here is to find the period r of a periodic function $f : B^n \rightarrow B^n$, i.e., a Boolean function for which

$$(244) f(x_i) = f(x_j) \text{ if and only if } x_j = x_i \oplus r, \text{ for all } x_i, x_j \in B^n$$

Note that since $x \oplus r \oplus r = x$, the function is 2-to-1.

Simon's algorithm solves the problem efficiently, with an exponential speed-up over any classical algorithm (see [Simon, 1994; Simon, 1997]). The algorithm proceeds as in the Deutsch-Jozsa algorithm, starting with the input and output registers in the state $|0 \dots 0\rangle|0\rangle$ in the computational basis:

$$(245) |0 \dots 0\rangle|0\rangle \xrightarrow{H} \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle|0\rangle$$

$$(246) \xrightarrow{U_f} \frac{1}{\sqrt{2^n}} \sum_x |x\rangle|f(x)\rangle$$

$$(247) = \frac{1}{\sqrt{2^{n-1}}} \sum_{x_i} \frac{|x_i\rangle + |x_i \oplus r\rangle}{\sqrt{2}} |f(x_i)\rangle$$

where U_f is the unitary transformation implementing the Boolean function as:

$$(248) U_f : |x\rangle|y\rangle \rightarrow |x\rangle|y \oplus f(x)\rangle$$

The usual way to see how the algorithm works is to consider what happens if we measure the output register and keep the state of the input register,¹⁶ which will have the form:

$$(249) \frac{|x_i\rangle + |x_i \oplus r\rangle}{\sqrt{2}}$$

This state contains the information r , but summed with an unwanted randomly chosen offset x_i that depends on the measurement outcome. A direct measurement of the state label would yield any $x \in B^n$ equiprobably, providing no information about r .

¹⁶The measurement of the output register here is a pedagogical device for ease of conceptualization. Only the input register is actually measured. The input register is in a mixture of states, which we can think of as the mixture associated with the distribution of outcomes obtained by measuring the output register.

We now apply a Hadamard transform:

$$(250) \quad \frac{|x_i\rangle + |x_i \oplus r\rangle}{\sqrt{2}} \xrightarrow{H} \frac{1}{\sqrt{2^n}} \sum_{y \in B^n} \frac{(-1)^{x_i \cdot y} + (-1)^{(x_i \oplus r) \cdot y}}{\sqrt{2}} |y\rangle$$

$$(251) \quad = \sum_{y:r \cdot y=0} \frac{(-1)^{x_i \cdot y}}{\sqrt{2^{n-1}}} |y\rangle$$

where the last equality follows because terms interfere destructively if $r \cdot y = 1$. Finally, we measure the input register in the computational basis and obtain a value y (equiprobably) such that $r \cdot y = 0$. Then we repeat the algorithm sufficiently many times to find enough values y_i so that r can be determined by solving the linear equations $r \cdot y_1 = 0, \dots, r \cdot y_k = 0$.

To see what is going on geometrically, consider the case $n = 2$. The possible values of the period r are: 01, 10, 11, and the corresponding states of the input and output registers after the unitary transformation U_f are:

$$r = 01 : \frac{1}{2}(|00\rangle + |01\rangle)|f(00)\rangle + \frac{1}{2}(|10\rangle + |11\rangle)|f(10)\rangle$$

$$r = 10 : \frac{1}{2}(|00\rangle + |10\rangle)|f(00)\rangle + \frac{1}{2}(|01\rangle + |11\rangle)|f(01)\rangle$$

$$r = 11 : \frac{1}{2}(|00\rangle + |11\rangle)|f(00)\rangle + \frac{1}{2}(|01\rangle + |10\rangle)|f(01)\rangle$$

Notice that this case reduces to the same geometric construction as in Deutsch's XOR algorithm. For $r = 10$ the input register states are $|c_1\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |10\rangle)$ or $|c_2\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |11\rangle)$, and for $r = 11$ the input register states are $|b_1\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ or $|b_2\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$, depending on the outcome of the measurement of the output register. So the three possible periods are associated with three planes in $\mathcal{H}^2 \otimes \mathcal{H}^2$, which correspond to the constant and balanced planes in Deutsch's XOR algorithm, and a third plane, all three planes intersecting in the line spanned by the vector $|00\rangle$. In the prime basis obtained by applying the Hadamard transformation, the planes are as follows:

$$r = 01 : \text{plane spanned by } |0'0'\rangle, |1'0'\rangle$$

$$r = 10 : \text{plane spanned by } |0'0'\rangle, |0'1'\rangle \text{ (corresponds to 'constant' plane)}$$

$$r = 11 : \text{plane spanned by } |0'0'\rangle, |1'1'\rangle \text{ (corresponds to 'balanced' plane)}$$

We could simply measure the input register in the prime basis to find the period. Alternatively, we could apply a Hadamard transformation (which amounts to dropping the primes in the above representation of the r -planes) and measure in the computational basis. The three planes are orthogonal, except for their intersection in the line spanned by the vector $|00\rangle$. The three possible periods can therefore be distinguished by measuring the observable with eigenstates $|00\rangle, |01\rangle, |10\rangle, |11\rangle$, except when the state of the register is projected by the measurement ('collapses')

onto the intersection state $|00\rangle$ (which occurs with probability $1/2$). So the algorithm will generally have to be repeated until we find an outcome that is not 00 .

The $n = 2$ case of Simon's algorithm reduces to Deutsch's XOR algorithm. What about other cases? We can see what happens in the general case if we consider the case $n = 3$. There are now seven possible periods: $001, 010, 011, 100, 101, 110, 111$. For the period $r = 001$, the state of the two registers after the unitary transformation U_f is:

$$\begin{aligned} & \frac{1}{2\sqrt{2}}(|000\rangle + |001\rangle)|f(000)\rangle + \frac{1}{2\sqrt{2}}(|010\rangle + |011\rangle)|f(010)\rangle \\ & + \frac{1}{2\sqrt{2}}(|100\rangle + |101\rangle)|f(100)\rangle + \frac{1}{2\sqrt{2}}(|110\rangle + |111\rangle)|f(110)\rangle \end{aligned}$$

If we measure the output register, the input register is left in one of four states, depending on the outcome of the measurement:

$$\begin{aligned} \frac{1}{\sqrt{2}}(|000\rangle + |001\rangle) &= \frac{1}{2}(|0'0'0'\rangle + |0'1'0'\rangle + |1'0'0'\rangle + |1'1'0'\rangle) \\ \frac{1}{\sqrt{2}}(|010\rangle + |011\rangle) &= \frac{1}{2}(|0'0'0'\rangle - |0'1'0'\rangle + |1'0'0'\rangle - |1'1'0'\rangle) \\ \frac{1}{\sqrt{2}}(|100\rangle + |101\rangle) &= \frac{1}{2}(|0'0'0'\rangle + |0'1'0'\rangle - |1'0'0'\rangle - |1'1'0'\rangle) \\ \frac{1}{\sqrt{2}}(|110\rangle + |111\rangle) &= \frac{1}{2}(|0'0'0'\rangle - |0'1'0'\rangle - |1'0'0'\rangle + |1'1'0'\rangle) \end{aligned}$$

Applying a Hadamard transformation amounts to dropping the primes. So if the period is $r = 001$, the state of the input register ends up in the 4-dimensional subspace of $\mathcal{H}^2 \otimes \mathcal{H}^2 \otimes \mathcal{H}^2$ spanned by the vectors: $|000\rangle, |010\rangle, |100\rangle, |110\rangle$.

A similar analysis applies to the other six possible periods. The corresponding subspaces are spanned by the following vectors:

$$r = 001: |000\rangle, |010\rangle, |100\rangle, |110\rangle$$

$$r = 010: |000\rangle, |001\rangle, |100\rangle, |101\rangle$$

$$r = 011: |000\rangle, |011\rangle, |100\rangle, |111\rangle$$

$$r = 100: |000\rangle, |001\rangle, |010\rangle, |011\rangle$$

$$r = 101: |000\rangle, |010\rangle, |101\rangle, |111\rangle$$

$$r = 110: |000\rangle, |001\rangle, |110\rangle, |111\rangle$$

$$r = 111: |000\rangle, |011\rangle, |101\rangle, |110\rangle$$

These subspaces are orthogonal except for intersections in 2-dimensional planes. The period can be found by measuring in the computational basis. Repetitions of the measurement will eventually yield sufficiently many distinct values to determine in which subspace out of the seven possibilities the final state lies. In this case ($n = 3$), it is clear by examining the above list that two values distinct from 000 suffice to determine the subspace, and these are just the values y_i for which $y_i \cdot r = 0$. Note that the subspaces correspond to quantum disjunctions. So determining the period of the function by Simon's algorithm amounts to determining which disjunction out of the seven alternative disjunctions is true, i.e., which subspace contains the state, without determining the truth values of the disjuncts.

Shor's Algorithm

Shor's factorization algorithm exploits the fact that the two prime factors p, q of a positive integer $N = pq$ can be found by determining the period of a function $f(x) = a^x \bmod N$, for any $a < N$ which is coprime to N , i.e., has no common factors with N other than 1. The period r of $f(x)$ depends on a and N . Once we know the period, we can factor N if r is even and $a^{r/2} \not\equiv -1 \pmod N$, which will be the case with probability greater than $1/2$ if a is chosen randomly. (If not, we choose another value of a .) The factors of N are the greatest common factors of $a^{r/2} \pm 1$ and N , which can be found in polynomial time by the Euclidean algorithm. (For these number-theoretic results, see [Nielsen and Chuang, 2000, Appendix 4].) So the problem of factorizing a composite integer N that is the product of two primes reduces to the problem of finding the period of a certain periodic function $f : Z_s \rightarrow Z_N$, where Z_n is the additive group of integers mod n (rather than B^n , the n -fold Cartesian product of a Boolean algebra B , as in Simon's algorithm). Note that $f(x+r) = f(x)$ if $x+r \leq s$. The function f is periodic if r divides s exactly, otherwise it is almost periodic.

Consider first the general form of the algorithm, as it is usually formulated. We begin by initializing the input register (s qubits) to the state $|0\rangle \in \mathcal{H}^s$ and the output register (N qubits) to the state $|0\rangle \in \mathcal{H}^N$. An s -fold Hadamard transformation is applied to the input register, followed by the unitary transformation U_f which implements the function $f(x) = a^x \bmod N$:

$$(252) \quad |0\rangle|0\rangle \xrightarrow{H} \frac{1}{\sqrt{s}} \sum_{x=0}^{s-1} |x\rangle|0\rangle$$

$$(253) \quad \xrightarrow{U_f} \frac{1}{\sqrt{s}} \sum_{x=0}^{s-1} |x\rangle|x + a^x \bmod N\rangle$$

Then we measure the output register in the computational basis¹⁷ and obtain a

¹⁷As in the discussion of Simon's algorithm, this measurement is purely hypothetical, introduced to simplify the analysis. Only the input register is actually measured.

state of the following form for the input register:

$$(254) \quad \frac{1}{\sqrt{s/r}} \sum_{j=0}^{s/r-1} |x_i + jr\rangle$$

This will be the case if r divides s exactly. The value x_i is the offset, which depends on the outcome i of the measurement of the output register. The sum is taken over the values of j for which $f(x_i + jr) = i$. When r does not divide s exactly, the analysis is a little more complicated. For a discussion, see [Barenco, 1998, 164], [Jozsa, 1997b]. Since the state label contains the random offset, a direct measurement of the label yields no information about the period.

A discrete Fourier transform for the integers mod s is now applied to the input register, i.e., a unitary transformation:

$$(255) \quad |x\rangle \xrightarrow{U_{DFT_s}} \frac{1}{\sqrt{s}} \sum_{y=0}^{s-1} e^{2\pi i \frac{xy}{s}} |y\rangle, \text{ for } x \in Z_s$$

This yields the transition:

$$(256) \quad \frac{1}{\sqrt{\frac{s}{r}}} \sum_{j=0}^{\frac{s}{r}-1} |x_i + jr\rangle \xrightarrow{U_{DFT_s}} \frac{1}{\sqrt{r}} \sum_{k=0}^{r-1} e^{2\pi i \frac{x_i k}{r}} |ks/r\rangle$$

and so shifts the offset into a phase factor and inverts the period as a multiple of s/r . A measurement of the input register in the computational basis yields $c = ks/r$. The algorithm is run a number of times until a value of k coprime to r is obtained. Cancelling c/s to lowest terms then yields k and r as k/r .

Since we don't know the value of r in advance of applying the algorithm, we do not, of course, recognize when a measurement outcome yields a value of k coprime to r . The idea is to run the algorithm, cancel c/s to lowest terms to obtain a candidate value for r and hence a candidate factor of N , which can then be tested by division into N . Even when we do obtain a value of k coprime to r , some values of a will yield a period for which the method fails to yield a factor of N , in which case we randomly choose a new value of a and run the algorithm with this value. The point is that all these steps are efficient, i.e., can be performed in polynomial time, and since only a polynomial number of repetitions are required to determine a factor with any given probability $p < 1$, the algorithm is a polynomial-time algorithm, achieving an exponential speed-up over any known classical algorithm.

To see how the algorithm functions geometrically, consider the case $N = 15$, $a = 7$ and $s = 64$ discussed in [Barenco, 1998, 160]. In this case, the function $f(x) =$

$a^x \bmod 15$ is:

$$\begin{aligned} 7^0 \bmod 15 &= 1 \\ 7^1 \bmod 15 &= 7 \\ 7^2 \bmod 15 &= 4 \\ 7^3 \bmod 15 &= 13 \\ 7^4 \bmod 15 &= 1 \\ &\vdots \\ 7^{63} \bmod 15 &= 13 \end{aligned}$$

and the period is evidently $r = 4$.¹⁸ After the application of the unitary transformation $U_f = a^x \bmod N$, the state of the two registers is:

$$\begin{aligned} &\frac{1}{8}(|0\rangle|1\rangle + |1\rangle|7\rangle + |2\rangle|4\rangle + |3\rangle|13\rangle \\ &\quad + |4\rangle|1\rangle + |5\rangle|7\rangle + |6\rangle|4\rangle + |7\rangle|13\rangle \\ &\quad \quad \quad \vdots \\ (257) \quad &+ |60\rangle|1\rangle + |61\rangle|7\rangle + |62\rangle|4\rangle + |63\rangle|13\rangle \end{aligned}$$

This is the state (253) for $s = 64$, $a = 7$. This state can be expressed as:

$$\begin{aligned} &\frac{1}{4}(|0\rangle + |4\rangle + |8\rangle + \dots + |60\rangle)|1\rangle \\ &\quad + \frac{1}{4}(|1\rangle + |5\rangle + |9\rangle + \dots + |61\rangle)|7\rangle \\ &\quad + \frac{1}{4}(|2\rangle + |6\rangle + |10\rangle + \dots + |62\rangle)|4\rangle \\ (258) \quad &+ \frac{1}{4}(|3\rangle + |7\rangle + |11\rangle + \dots + |63\rangle)|13\rangle \end{aligned}$$

If we measure the output register, we obtain (equiprobably) one of four states for the input register, depending on the outcome of the measurement: 1, 7, 4, or 13:

$$(259) \quad \frac{1}{4}(|0\rangle + |4\rangle + |8\rangle + \dots + |60\rangle)$$

$$(260) \quad \frac{1}{4}(|1\rangle + |5\rangle + |9\rangle + \dots + |61\rangle)$$

$$(261) \quad \frac{1}{4}(|2\rangle + |6\rangle + |10\rangle + \dots + |62\rangle)$$

$$(262) \quad \frac{1}{4}(|3\rangle + |7\rangle + |11\rangle + \dots + |63\rangle)$$

These are the states (254) for values of the offset $x_1 = 0$, $x_7 = 1$, $x_4 = 2$, $x_{13} = 3$. Application of the discrete Fourier transform yields:

$$\begin{aligned} x_1 = 0 &: \frac{1}{2}(|0\rangle + |16\rangle + |32\rangle + |48\rangle) \\ x_7 = 1 &: \frac{1}{2}(|0\rangle + i|16\rangle - |32\rangle - i|48\rangle) \\ x_4 = 2 &: \frac{1}{2}(|0\rangle - |16\rangle + |32\rangle - |48\rangle) \\ x_{13} = 3 &: \frac{1}{2}(|0\rangle - i|16\rangle - |32\rangle + i|48\rangle) \end{aligned}$$

¹⁸The factors 3 and 5 of 15 are derived as the greatest common factors of $a^{r/2} - 1 = 48$ and 15 and $a^{r/2} + 1 = 50$ and 15, respectively.

which are the states in (256). (Here $s = 64$, $r = 4$; $\sqrt{\frac{s}{r}} = 4$, $\frac{s}{r} - 1 = 15$.) So for the period $r = 4$, the state of the input register ends up in the 4-dimensional subspace spanned by the vectors $|0\rangle, |16\rangle, |32\rangle, |48\rangle$.

Now consider all possible even periods r for which $f(x) = a^x \bmod 15$, where a is coprime to 15. The other possible periods values of a are 2, 4, 8, 11, 13, 14 and the corresponding periods turn out to be 4, 2, 4, 2, 4, 2. So we need only consider $r = 2$.¹⁹Note that different values of a with the same period affect only the labels of the output register (e.g., for $a = 2$, the labels are $|1\rangle, |2\rangle, |4\rangle, |8\rangle$ instead of $|1\rangle, |7\rangle, |4\rangle, |13\rangle$).

For $r = 2$, if we measure the output register, we will obtain (equiprobably) one of two states for the input register, depending on the outcome of the measurement (say, a or b):

$$(263) \quad |0\rangle + |2\rangle + |4\rangle + \dots + |62\rangle$$

$$(264) \quad |1\rangle + |3\rangle + |5\rangle + \dots + |63\rangle$$

After the discrete Fourier transform, these states are transformed to:

$$x_a = 0: \quad |0\rangle + |32\rangle$$

$$x_b = 1: \quad |0\rangle - |32\rangle$$

In this case, the 2-dimensional subspace $\mathcal{V}_{r=2}$ spanned by $|0\rangle, |32\rangle$ for $r = 2$ is included in the 4-dimensional subspace $\mathcal{V}_{r=4}$ for $r = 4$. A measurement can distinguish $r = 4$ from $r = 2$ reliably, i.e., whether the final state of the input register is in $\mathcal{V}_{r=4}$ or $\mathcal{V}_{r=2}$, only if the final state is in $\mathcal{V}_{r=4} - \mathcal{V}_{r=2}$, the part of $\mathcal{V}_{r=4}$ orthogonal to $\mathcal{V}_{r=2}$. What happens if the final state ends up in $\mathcal{V}_{r=2}$?

Shor's algorithm works as a randomized algorithm. As mentioned above, it produces a candidate value for the period r and hence a candidate factor of N , which can be tested (in polynomial time) by division into N . A measurement of the input register in the computational basis yields an outcome $c = ks/r$. The value of k is chosen equiprobably by the measurement of the output register. The procedure is to repeat the algorithm until the outcome yields a value of k coprime to r , in which case canceling c/s to lowest terms yields k and r as k/r .

For example, suppose we choose $a = 7$, in which case (unknown to us) $r = 4$. The values of k coprime to r are $k = 1$ and $k = 3$ (this is also unknown to us, because k depends on the value of r). Then c/s cancelled to lowest terms is $1/4$ and $3/4$, respectively, both of which yield the correct period. From the geometrical perspective, these values of k correspond to finding the state after measurement in the computational basis to be $|16\rangle$ or $|48\rangle$, both of which do distinguish $\mathcal{V}_{r=4}$ from $\mathcal{V}_{r=2}$.

Suppose we choose a value of a with period $r = 2$ and find the value $c = 32$. The only value of k coprime to r is $k = 1$. Then c/s cancelled to lowest terms is $1/2$, which yields the correct period, and hence the correct factors of N . But $c = 32$

¹⁹Every value of a except $a = 14$ yields the correct factors for 15. For $a = 14$, the method fails: $r = 2$, so $a^{\frac{r}{2}} = -1 \bmod 15$.

could also be obtained for $a = 7$, $r = 4$, and $k = 2$, which does not yield the correct period, and hence does not yield the correct factors of N . Putting it geometrically: the value $k = 1$ for $r = 2$ corresponds to the same state, $|32\rangle$, as the value $k = 2$ for $r = 4$. Once we obtain the candidate period $r = 2$ (by cancelling $c/s = 32/64$ to lowest terms), we calculate the factors of N as the greatest common factors of $a \pm 1$ and N and test these by division into N . If $a = 7$, these calculated factors will be incorrect. If $a = 2$, say, the factors calculated in this way will be correct.

We see that, with the added information provided by the outcome of a test division of a candidate factor into N , Shor's randomized algorithm again amounts to determining which disjunction among alternative disjunctions is true, i.e., which subspace contains the state, without determining the truth values of the disjuncts.

6.3 Where Does the Speed-Up Come From?

What, precisely, is the feature of a quantum computer responsible for the phenomenal efficiency over a classical computer? In the case of Simon's algorithm, the speed-up is exponential over any classical algorithm; in the case of Shor's algorithm, the speed-up is exponential over any known classical algorithm.

Steane [1998] remarks:

The period finding algorithm appears at first sight like a conjuring trick: it is not quite clear how the quantum computer managed to produce the period like a rabbit out of a hat. ... I would say that the most important features are contained in $[|\psi\rangle = \frac{1}{s} \sum_{x=0}^{s-1} |x\rangle |f(x)\rangle]$. They are not only the *quantum parallelism* already mentioned, but also *quantum entanglement*, and, finally, quantum interference. Each value of $f(x)$ retains a link with the value of x which produced it, through entanglement of the x and y registers in $[|\psi\rangle]$. The 'magic' happens when a measurement of the y register produces the special state $[\frac{1}{s/r} \sum_{j=0}^{s/r-1} |x_i + jr\rangle]$ in the x -register, and it is quantum entanglement which permits this (see also [Jozsa, 1997a]). The final Fourier transform can be regarded as an interference between the various superposed states in the x -register (compare with the action of a diffraction grating).

Interference effects can be used for computational purposes with classical light fields, or water waves for that matter, so interference is not in itself the essentially quantum feature. Rather, the exponentially large number of interfering states, and the entanglement, are features which do not arise in classical systems.

Jozsa points out [1997a] that the state space (phase space) of a composite classical system is the Cartesian product of the state spaces of its subsystems, while the state space of a composite quantum system is the tensor product of the state spaces of its subsystems. For n qubits, the quantum state space has 2^n dimensions. So the information required to represent a general state increases exponentially with n : even if we restrict the specification of the amplitudes to numbers of finite precision, a superposition will in general have $\mathcal{O}(2^n)$ components. For a classical composite system of n two-level subsystems, the number of possible states grows exponentially with n , but the information required to represent a general state is just n times the information required to represent a single two-level system, i.e., the information grows only linearly with n because the state of a composite system is just a product state.

More formally, Jozsa and Linden [2002] have shown that a quantum algorithm operating on pure states can achieve an exponential speed-up over classical algorithms only if the quantum algorithm involves multi-partite entanglement that increases unboundedly with the input size. Similarly, Vidal [2003] has shown that a classical computer can simulate the evolution of a pure state of n qubits with computational resources that grow linearly with n and exponentially in multi-partite entanglement.

The essential feature of the quantum computations discussed above in §6.2 is the selection of a disjunction, representing a global property of a function, among alternative possible disjunctions without computing the truth values of the disjuncts, which is redundant information in a quantum computation but essential information classically. Note that a quantum disjunction is represented by a subspace of entangled states in the tensor product Hilbert space of the input and output registers. This is analogous to the procedure involved in the key observation underlying the proof of the quantum bit commitment theorem discussed in §5.2. The series of operations described by equations (167)–(172), in which the channel particle is entangled with ancilla systems and the ancillas are subsequently measured, effectively constitute a quantum computation.

The first stage of a quantum algorithm involves the creation of a state in which every input value to the function is correlated with a corresponding output value. This is referred to as ‘quantum parallelism’ and is sometimes cited as the source of the speed-up in a quantum computation. The idea is that a quantum computation is something like a massively parallel classical computation, for all possible values of a function. This appears to be Deutsch’s view, with the parallel computations taking place in parallel universes. For a critique, see [Steane, 2003], who defends a view similar to that presented here. Of course, all these different values are inaccessible: a measurement in the computational basis will only yield (randomly) one correlated input-output pair. Further processing is required, including the final discrete Fourier transform for the three algorithms discussed in §6.2. It would be incorrect to attribute the efficiency of these quantum algorithms to the interference in the input register produced by the Fourier transform. The role of the Fourier transform is simply to allow a measurement in the computational basis to reveal which subspace representing the target disjunction contains the state.

One might wonder, then, why the discrete Fourier transform is even necessary. We could, of course, simply perform an equivalent measurement in a different basis. But note that a computation would have to be performed to determine this basis. This raises the question of precisely how to assess the speed-up of a quantum algorithm relative to a rival classical algorithm. What are the relevant computational steps to be counted in making this assessment for a quantum computation? Since any sequence of unitary transformations is equivalent to a single unitary transformation, and a unitary transformation followed by a measurement in a certain basis is equivalent to simply performing a measurement in a different basis, any quantum computation can always be reduced to just one step: a measurement in a particular basis!

Of course, this observation is hardly illuminating, since a computation at least as difficult as the original computation would have to be performed to determine the required basis, but it does indicate that some convention is required about what steps to count in a quantum computation. The accepted convention is to require the unitary transformations in a quantum computation to be constructed from elementary quantum gates that form a universal set (e.g., the CNOT gate, the Hadamard gate, the phase gate, and the $\pi/8$ gate discussed in §6.1) and to count each such gate as one step. In addition, all measurements are required to be performed in the computational basis, and these are counted as additional steps. The final discrete Fourier transforms in the Deutsch-Jozsa algorithm, Simon's algorithm, and Shor's algorithm are indispensable in transforming the state so that the algorithms can be completed by measurements in the computational basis, and it is an important feature of these algorithms that the Fourier transform can be implemented efficiently with elementary unitary gates. To claim that a quantum algorithm is exponentially faster than a classical algorithm is to claim that the number of steps counted in this way for the quantum algorithm is a polynomial function of the size of the input (the number of qubits required to store the input), while the classical algorithm involves a number of steps that increases exponentially with the size of the input (the number of bits required to store the input).

7 QUANTUM FOUNDATIONS FROM THE PERSPECTIVE OF QUANTUM INFORMATION

Does the extension of the classical theory of information to quantum states shed new light on the foundational problems of quantum mechanics underlying the Bohr-Einstein debate mentioned in §1, in particular the measurement problem? Researchers in the area of quantum information and quantum computation often suggest a positive answer to this question, with a promissory note for how the story is supposed to go. More fully worked-out (generally, rather different) positive responses have been proposed by various authors, notably Fuchs [2001b; 2002a; 2002b; 2001a] and Brukner and Zeilinger [2001; 20002]. For a very thorough analysis and critique of the Brukner-Zeilinger position, see [Timpson, 2004]. See also Hall [2000] and the response by Brukner and Zeilinger [2000]. Here I shall limit my discussion to the significance of a characterization theorem for quantum mechanics in terms of information-theoretic constraints by Clifton, Bub, and Halvorson (CBH) [2003].

7.1 The CBH Characterization Theorem

CBH showed that one can derive the basic kinematic features of a quantum description of physical systems from three fundamental information-theoretic constraints:

- the impossibility of superluminal information transfer between two physical

systems by performing measurements on one of them,

- the impossibility of perfectly broadcasting the information contained in an unknown physical state (which, for pure states, amounts to ‘no cloning’),
- the impossibility of communicating information so as to implement a bit commitment protocol with unconditional security (so that cheating is in principle excluded by the theory).

More precisely, CBH formulate these information-theoretic constraints in the general framework of C^* -algebras, which allows a mathematically abstract characterization of a physical theory that includes, as special cases, all classical mechanical theories of both wave and particle varieties, and all variations on quantum theory, including quantum field theories (plus any hybrids of these theories, such as theories with superselection rules). Within this framework, CBH show that the three information-theoretic constraints jointly entail three physical conditions that they take as definitive of what it means to be a quantum theory in the most general sense. Specifically, the information-theoretic constraints entail that:

- the algebras of observables pertaining to distinct physical systems commute (a condition usually called microcausality or, to use Summers’ term [Summers, 1990], *kinematic independence*),
- any individual system’s algebra of observables is *noncommutative*,
- the physical world is *nonlocal*, in that spacelike separated systems can occupy entangled states that persist as the systems separate.

CBH also partly demonstrated the converse derivation, leaving open a question concerning nonlocality and bit commitment. This remaining issue was later resolved by Hans Halvorson [2004], so the CBH theorem is a characterization theorem for quantum theory in terms of the three information-theoretic constraints.

Note that the C^* -algebraic framework is not restricted to the standard quantum mechanics of a system represented on a single Hilbert space with a unitary dynamics, but is general enough to cover cases of systems with an infinite number of degrees of freedom that arise in quantum field theory and the thermodynamic limit of quantum statistical mechanics (in which the number of microsystems and the volume they occupy goes to infinity, while the density defined by their ratio remains constant). The C^* -algebraic framework has even been applied to quantum field theory on curved spacetime and so is applicable to the quantum theoretical description of exotic phenomena such as Hawking radiation (black hole evaporation); see [Wald, 1984]. The Stone-von Neumann theorem, which guarantees the existence of a unique representation (up to unitary equivalence) of the canonical commutation relations for systems with a finite number of degrees of freedom, breaks down for such cases, and there will be many unitarily inequivalent representations of the canonical commutation relations.

One could, of course, consider weaker mathematical structures, but it seems that the C^* -algebraic machinery suffices for all physical theories that have been found to be empirically successful to date, including phase space theories and Hilbert space theories [Landsman, 1998], and theories based on a manifold [Connes, 1994]. For further discussion of this point, see Halvorson and Bub [2005]. See also Halvorson (this vol., chap. 8), Emch (this vol., ch. 8), and Landsman (this vol., ch. 5).

A C^* -algebra is essentially an abstract generalization of the structure of the algebra of operators on a Hilbert space. Technically, a (unital) C^* -algebra is a Banach $*$ -algebra over the complex numbers containing the identity, where the involution operation $*$ and the norm are related by $\|A^*A\| = \|A\|^2$. So the algebra $\mathfrak{B}(\mathcal{H})$ of all bounded operators on a Hilbert space \mathcal{H} is a C^* -algebra, with $*$ the adjoint operation and $\|\cdot\|$ the standard operator norm.

In standard quantum theory, as discussed in §3.1, a state on $\mathfrak{B}(\mathcal{H})$ is defined by a density operator ρ on \mathcal{H} in terms of an expectation-valued functional $\tilde{\rho}(A) = \text{Tr}(\rho A)$ for all observables represented by self-adjoint operators A in $\mathfrak{B}(\mathcal{H})$. This definition of $\tilde{\rho}(A)$ in terms of ρ yields a positive normalized linear functional. So a state on a C^* -algebra \mathfrak{C} is defined, quite generally, as any positive normalized linear functional $\tilde{\rho} : \mathfrak{C} \rightarrow \mathbb{C}$ on the algebra. Pure states can be defined by the condition that if $\tilde{\rho} = \lambda\tilde{\rho}_1 + (1 - \lambda)\tilde{\rho}_2$ with $\lambda \in (0, 1)$, then $\tilde{\rho} = \tilde{\rho}_1 = \tilde{\rho}_2$; other states are mixed. In the following, we drop the ‘ \sim ’ in $\tilde{\rho}$, but note that a C^* -algebraic state ρ is a positive linear functional on \mathfrak{C} , while the density operator of standard quantum mechanics is an element of $\mathfrak{C} = \mathfrak{B}(\mathcal{H})$.

By Gleason’s theorem [Gleason, 1957], every C^* -algebraic state in this sense on a C^* -algebra $\mathfrak{C} = \mathfrak{B}(\mathcal{H})$ is given by a density operator on $\mathfrak{B}(\mathcal{H})$. However, because countable additivity is not presupposed by the C^* -algebraic notion of state (and, therefore, Gleason’s theorem does not apply in general), there can be pure states of $\mathfrak{B}(\mathcal{H})$ that are not representable by vectors in \mathcal{H} . In fact, if A is any self-adjoint element of a C^* -algebra \mathfrak{A} , and $a \in \text{sp}(A)$, then there always exists a pure state ρ of \mathfrak{A} that assigns a *dispersion-free* value of a to A [Kadison and Ringrose, 1997, Ex. 4.6.31]. Since this is true even when we consider a point in the continuous spectrum of a self-adjoint operator A acting on a Hilbert space, *without* any corresponding eigenvector, it follows that there *are* pure states of $\mathfrak{B}(\mathcal{H})$ in the C^* -algebraic sense that cannot be vector states (nor, in fact, representable by any density operator \mathcal{H}).

As we saw in §3.1, the general evolution of a quantum system resulting from a combination of unitary interactions and selective or nonselective measurements can be described by a quantum operation, i.e., a completely positive linear map. Accordingly, a completely positive linear map $T : \mathfrak{C} \rightarrow \mathfrak{C}$, where $0 \leq T(I) \leq I$ is taken as describing the general evolution of a system represented by a C^* -algebra of observables. The map or operation T is called selective if $T(I) < I$ and nonselective if $T(I) = I$. Recall that a yes-no measurement of some idempotent observable represented by a projection operator P is an example of a selective operation. Here, $T(A) = PAP$ for all A in the C^* -algebra \mathfrak{C} , and ρ^T , the transformed (‘collapsed’) state, is the final state obtained after measuring P in the state ρ and ignoring

all elements of the ensemble that do not yield the eigenvalue 1 of P (so $\rho^T(A) = \rho(T(A))/\rho(T(I))$ when $\rho(T(I)) \neq 0$, and $\rho^T = 0$ otherwise). The time evolution in the Heisenberg picture induced by a unitary operator $U \in \mathfrak{C}$ is an example of a nonselective operation. Here, $T(A) = UAU^{-1}$. Similarly, the measurement of an observable O with spectral measure $\{P_i\}$, without selecting a particular outcome, is an example of a nonselective operation, with $T(A) = \sum_{i=1}^n P_i A P_i$. As in the standard quantum theory of a system with a finite-dimensional Hilbert space (cf. Eq. (67) of §3.1), any completely positive linear map can be regarded as the restriction to a local system of a unitary map on a larger system.

A representation of a C^* -algebra \mathfrak{C} is any mapping $\pi : \mathfrak{C} \rightarrow \mathfrak{B}(\mathcal{H})$ that preserves the linear, product, and $*$ structure of \mathfrak{C} . The representation is faithful if π is one-to-one, in which case $\pi(\mathfrak{C})$ is an isomorphic copy of \mathfrak{C} . The Gelfand-Naimark theorem says that every abstract C^* -algebra has a concrete faithful representation as a norm-closed $*$ -subalgebra of $\mathfrak{B}(\mathcal{H})$, for some appropriate Hilbert space \mathcal{H} . As indicated above, in the case of systems with an infinite number of degrees of freedom (e.g., quantum field theory), there are inequivalent representations of the C^* -algebra of observables defined by the commutation relations.

Every classical phase space theory defines a commutative C^* -algebra. For example, the observables of a classical system of n particles — the real-valued continuous functions on the phase space \mathbb{R}^{6n} — can be represented as the self-adjoint elements of the C^* -algebra $\mathfrak{B}(\mathbb{R}^{6n})$ of all continuous complex-valued functions f on \mathbb{R}^{6n} . The phase space \mathbb{R}^{6n} is locally compact and can be made compact by adding just one point ‘at infinity’, or we can simply consider a bounded (and thus compact) subset of \mathbb{R}^{6n} . The statistical states of the system are given by probability measures μ on \mathbb{R}^{6n} , and pure states, corresponding to maximally complete information about the particles, are given by the individual points of \mathbb{R}^{6n} . The system’s state ρ in the C^* -algebraic sense is the expectation functional corresponding to μ , defined by $\rho(f) = \int_{\mathbb{R}^{6n}} f d\mu$. Conversely [Kadison and Ringrose, 1997, Thm. 4.4.3], every commutative C^* -algebra \mathfrak{C} is isomorphic to the set $C(X)$ of all continuous complex-valued functions on a locally compact Hausdorff space X defined by the pure states of \mathfrak{C} . If \mathfrak{C} has a multiplicative identity, the ‘phase space’ X is compact. In this ‘function representation’ of C , the isomorphism maps an element $C \in \mathfrak{C}$ to the function \hat{C} (the Gelfand transformation of C) whose value at any ρ is just the (dispersion-free) value that ρ assigns to C . So ‘behind’ every abstract commutative C^* -algebra there is a classical phase space theory defined by its function representation on the phase space X . This representation theorem (and its converse) justifies characterizing a C^* -algebraic theory as classical just in case its algebra is commutative.

As we saw above, CBH identify quantum theories with a certain subclass of noncommutative C^* -algebras, where the condition of kinematic independence is satisfied by the algebras of observables of distinct systems and the states of space-like separated systems are characterized by the sort of nonlocality associated with entanglement.

To clarify the rationale for this characterization and the significance of the information-theoretic constraints, consider a composite quantum system AB , consisting of two subsystems, A and B . For simplicity, assume the systems are indistinguishable, so their C^* -algebras \mathfrak{A} and \mathfrak{B} are isomorphic. The observables of the component systems A and B are represented by the self-adjoint elements of \mathfrak{A} and \mathfrak{B} , respectively. Let $\mathfrak{A} \vee \mathfrak{B}$ denote the C^* -algebra generated by \mathfrak{A} and \mathfrak{B} . The physical states of A , B , and AB , are given by positive normalized linear functionals on their respective algebras that encode the expectation values of all observables. To capture the idea that A and B are *physically distinct* systems, CBH make the assumption that any state of \mathfrak{A} is compatible with any state of \mathfrak{B} , i.e., for any state ρ_A of \mathfrak{A} and ρ_B of \mathfrak{B} , there is a state ρ of $\mathfrak{A} \vee \mathfrak{B}$ such that $\rho|_{\mathfrak{A}} = \rho_A$ and $\rho|_{\mathfrak{B}} = \rho_B$.

The sense of the ‘no superluminal information transfer via measurement’ constraint is that when Alice and Bob, say, perform local measurements, Alice’s measurements can have no influence on the statistics for the outcomes of Bob’s measurements, and conversely. That is, merely performing a local measurement cannot, in and of itself, convey any information to a physically distinct system, so that everything ‘looks the same’ to that system after the measurement operation as before, in terms of the expectation values for the outcomes of measurements. CBH show [2003, Thm. 1] that it follows from this constraint that A and B are *kinematically independent* systems if they are physically distinct in the above sense, i.e., every element of \mathfrak{A} commutes pairwise with every element of \mathfrak{B} . (More precisely, an operation T on $\mathfrak{A} \vee \mathfrak{B}$ conveys no information to Bob just in case $(T^*\rho)|_{\mathfrak{B}} = \rho|_{\mathfrak{B}}$ for all states ρ of \mathfrak{B} , where T^* is the map on the states, i.e., the positive linear functionals on $\mathfrak{A} \vee \mathfrak{B}$, induced by T . Clearly, the kinematic independence of \mathfrak{A} and \mathfrak{B} entails that Alice’s local measurement operations cannot convey any information to Bob, i.e., $T(B) = \sum_{i=1}^n E_i^{1/2} B E_i^{1/2} = B$ for $B \in \mathfrak{B}$ if T is implemented by a POVM in \mathfrak{A} . CBH prove that if Alice cannot convey any information to Bob by performing local measurement operations, then \mathfrak{A} and \mathfrak{B} are kinematically independent.)

The ‘no broadcasting’ condition now ensures that the individual algebras \mathfrak{A} and \mathfrak{B} are noncommutative. Recall that for pure states, broadcasting reduces to cloning, and that in elementary quantum mechanics, neither cloning nor broadcasting is possible in general (see section 3.2). CBH show that broadcasting and cloning are always possible for classical systems, i.e., in the commutative case there *is* a universal broadcasting map that clones any pair of input pure states and broadcasts any pair of input mixed states [Clifton *et al.*, 2003, Thm. 2]. Conversely, they show that if any two states can be (perfectly) broadcast, then any two pure states can be cloned; and if two pure states of a C^* -algebra can be cloned, then they must be orthogonal. So, if any two states can be broadcast, then all pure states are orthogonal, from which it follows that the algebra is commutative.

The quantum mechanical phenomenon of interference is the physical manifestation of the noncommutativity of quantum observables or, equivalently, the superposition of quantum states. So the impossibility of perfectly broadcasting the

information contained in an unknown physical state, or of cloning or copying the information in an unknown pure state, is the information-theoretic counterpart of interference.

Now, if \mathfrak{A} and \mathfrak{B} are noncommutative and mutually commuting, it can be shown that there are nonlocal entangled states on the C^* -algebra $\mathfrak{A} \vee \mathfrak{B}$ they generate (see [Landau, 1987; Summers, 1990; Bacciagaluppi, 1994], and — more relevantly here, in terms of a specification of the range of entangled states that can be guaranteed to exist — [Halvorson, 2004]). So it seems that entanglement — what Schrödinger [1935, 555] identified as ‘*the* characteristic trait of quantum mechanics, the one that enforces its entire departure from classical lines of thought’, as we saw in §4.1 — follows automatically in any theory with a noncommutative algebra of observables. That is, it seems that once we assume ‘no superluminal information transfer via measurement’, and ‘no broadcasting’, the class of allowable physical theories is restricted to those theories in which physical systems manifest both interference *and* nonlocal entanglement. But in terms of physical interpretation this conclusion is a bit too quick, since the derivation of entangled states depends on formal properties of the C^* -algebraic machinery. Moreover, we have no assurance that two systems in an entangled state will maintain their entanglement indefinitely as they separate in space, which is the case for quantum entanglement. But this is precisely what is required by the cheating strategy that thwarts secure bit commitment, since Alice will have to keep one system of such a pair and send the other system to Bob, whose degree of spatial separation from Alice is irrelevant, in principle, to the implementation of the protocol. In an information-theoretic characterization of quantum theory, the fact that entangled states of composite systems can be instantiated, and instantiated nonlocally so that the entanglement of composite system is maintained as the subsystems separate in space, should be shown to follow from some information-theoretic principle. The role of the ‘no bit commitment’ constraint is to guarantee the persistence of entanglement over distance, i.e., the existence of a certain class of nonlocal entangled states — hence it gives us nonlocality, not merely ‘holism.’

As shown in §5.2, unconditionally secure quantum bit commitment is impossible because a generalized version of the EPR cheating strategy can always be applied by introducing additional ancilla particles and enlarging the Hilbert space in a suitable way. That is, for a quantum mechanical system consisting of two (separated) subsystems represented by the C^* -algebra $\mathfrak{B}(\mathcal{H}_1) \otimes \mathfrak{B}(\mathcal{H}_2)$, any mixture of states on $\mathfrak{B}(\mathcal{H}_2)$ can be generated from a distance by performing an appropriate generalized measurement on the system represented by $\mathfrak{B}(\mathcal{H}_1)$, for an appropriate entangled state of the composite system. This is what Schrödinger called ‘remote steering’ and found so physically counterintuitive that he speculated [1936, 451] (wrongly, as it turned out) that experimental evidence would eventually show that this was simply an artefact of the theory, not instantiated in our world. He suggested that an entangled state of a composite system would almost instantaneously decay to a mixture as the component systems separated.²⁰ There would still be

²⁰A similar possibility was raised and rejected by Furry [1936].

correlations between the states of the component systems, but remote steering would no longer be possible.

It seems worth noticing that the [EPR] paradox could be avoided by a very simple assumption, namely if the situation after separating were described by [the entangled state $\Psi(x, y) = \sum_k a_k g_k(x) f_k(y)$], but with the additional statement that the knowledge of the *phase relations* between the complex constants a_k has been entirely lost in consequence of the process of separation. This would mean that not only the parts, but the whole system, would be in the situation of a mixture, not of a pure state. It would not preclude the possibility of determining the state of the first system by *suitable* measurements in the second one or *vice versa*. But it would utterly eliminate the experimenters influence on the state of that system which he does not touch.

Schrödinger regarded the phenomenon of interference associated with noncommutativity in quantum mechanics as unproblematic, because he saw this as reflecting the fact that particles are wavelike. But he did not believe that we live in a world in which physical systems can exist nonlocally in entangled states, because such states would allow Alice to steer Bob's system into any mixture of pure states compatible with Bob's reduced density operator and he did not expect that experiments would bear this out. Of course, it was an experimental question in 1935 whether Schrödinger's conjecture was correct or not. We now know that the conjecture is false. A wealth of experimental evidence, including the confirmed violations of Bell's inequality [Aspect *et al.*, 1981; Aspect *et al.*, 1982] and the confirmations of quantum teleportation [Bouwmeester *et al.*, 1997; Boschi *et al.*, 1998; Furusawa *et al.*, 1998; Nielsen *et al.*, 1998], testify to this. The relevance of Schrödinger's conjecture here is this: it raises the possibility of a quantum-like world in which there is interference but no nonlocal entanglement. Can we exclude this possibility on information-theoretic grounds?

Now although unconditionally secure bit commitment is no less impossible for classical systems, in which the algebras of observables are commutative, than for quantum systems, the insecurity of a bit commitment protocol in a noncommutative setting depends on considerations entirely different from those in a classical commutative setting. As we saw in §5.2, the security of a classical bit commitment protocol is a matter of computational complexity and cannot be unconditional.

By contrast, if, as Schrödinger speculated, we lived in a world in which the algebras of observables are noncommutative but composite physical systems cannot exist in nonlocal entangled states, if Alice sends Bob one of two mixtures associated with the same density operator to establish her commitment, then she is, in effect, sending Bob evidence for the truth of an exclusive disjunction that is not based on the selection of a particular disjunct. (Bob's reduced density operator is associated ambiguously with both mixtures, and hence with the truth of the exclusive disjunction: '0 or 1'.) Noncommutativity allows the possibility of different mixtures associated with the same density operator. What thwarts the possibility

of using the ambiguity of mixtures in this way to implement an unconditionally secure bit commitment protocol is the existence of nonlocal entangled states between Alice and Bob. This allows Alice to cheat by preparing a suitable entangled state instead of one of the mixtures, where the reduced density operator for Bob is the same as that of the mixture. Alice is then able to steer Bob's systems remotely into either of the two mixtures associated with the alternative commitments at will.

So what *would* allow unconditionally secure bit commitment in a noncommutative theory is the absence of physically occupied nonlocal entangled states, or the spontaneous destruction of entanglement as systems separate. One can therefore take Schrödinger's remarks as relevant to the question of whether or not secure bit commitment is possible in our world. In effect, Schrödinger raised the possibility that we live in a quantum-like world in which unconditionally secure bit commitment is possible! It follows that the impossibility of unconditionally secure bit commitment entails that, for any mixed state that Alice and Bob can prepare by following some (bit commitment) protocol, there is a corresponding nonlocal entangled state that can be physically occupied by Alice's and Bob's particles and persists indefinitely as the particles move apart.

To sum up: the content of the CBH theorem is that a quantum theory — a C^* -algebraic theory whose observables and states satisfy conditions of kinematic independence, noncommutativity, and nonlocality — can be characterized by three information-theoretic constraints: no superluminal communication of information via measurement, no (perfect) broadcasting, and no (unconditionally secure) bit commitment.

7.2 *Quantum Mechanics as a Theory of Information*

Consider Einstein's view,²¹ mentioned in §1, that quantum mechanics is incomplete. Essentially, Einstein based his argument for this claim on the demand that a complete physical theory should satisfy certain principles of realism (essentially, a locality principle and a separability principle), which amounts to the demand that statistical correlations between spatially separated systems should have a common causal explanation in terms of causal factors obtaining at the common origin of the systems. Roughly thirty years after the publication of the Einstein-Podolsky-Rosen paper [1935], John Bell [1964] showed that the statistical correlations of the entangled Einstein-Podolsky-Rosen state for spatially separated particles are inconsistent with any explanation in terms of a classical probability distribution over common causal factors originating at the source of the particles before they separate. But the fact that quantum mechanics allows the possibility of correlations that are not reducible to common causes is a *virtue* of the theory. It is precisely the nonclassical correlations of entangled states that underlie the possibility of an exponential speed-up of quantum computation over classical computation, the

²¹The following discussion is adapted from [Bub, 2004] and [Bub, 2005], but the argument here is developed somewhat differently.

possibility of unconditionally secure key distribution but the impossibility of unconditionally secure quantum bit commitment, and phenomena such as quantum teleportation and other nonclassical entanglement-assisted communication protocols.

While Einstein's argument for incompleteness fails, there is another sense, also associated with entangled states, in which quantum mechanics might be said to be incomplete. In a typical (idealized) quantum mechanical measurement interaction, say an interaction in which the two possible values, 0 and 1, of an observable of a qubit in a certain quantum state become correlated with the two possible positions of a macroscopic pointer observable, p_0 and p_1 , the final state is an entangled state, a linear superposition of the states $|0\rangle|p_0\rangle$ and $|1\rangle|p_1\rangle$ with coefficients derived from the initial quantum state of the qubit. To dramatize the problem, Schrödinger [1936] considered the case where $|p_0\rangle$ and $|p_1\rangle$ represent the states of a cat being alive and a cat being dead in a closed box, which is only opened by the observer some time after the measurement interaction. On the standard way of relating the quantum state of a system to what propositions about the system are determinately (definitely) true or false, and what propositions have no determinate truth value, some correlational proposition about the composite system (microsystem + cat) is true in this entangled state, but the propositions asserting that the cat is alive (and the value of the qubit observable is 0), or that the cat is dead (and the value of the qubit observable is 1), are assigned no determinate truth value. Moreover, if we assume that the quantum propositions form an algebraic structure isomorphic to the structure of subspaces of the Hilbert space of the composite system — the representational space for quantum states and observables — then it is easy to derive a formal contradiction from the assumption that the correlational proposition corresponding to the entangled state is true, and that the cat is either definitely alive or definitely dead.²² Schrödinger thought that it was absurd to suppose that quantum mechanics requires us to say that the cat in such a situation (a macrosystem) is neither alive nor dead (does not have a determinate macroproperty of this sort) until an observer opens the box and looks, in which case the entangled state 'collapses' nonlinearly and stochastically, with probabilities given by the initial quantum state of the microsystem, onto a product of terms representing a definite state of the cat and a definite state of the microsystem. Einstein [1967, 39] concurred and remarked in a letter to Schrödinger: 'If that were so then physics could only claim the interest of shopkeepers and engineers; the whole thing would be a wretched bungle.'

This is the standard 'measurement problem' of quantum mechanics. Admittedly, the formulation of the problem is highly idealized, but the fundamental problem arises from the way in which quantum mechanics represents correlations via entangled states and does not disappear entirely in less idealized formulations (even though the problem is somewhat altered by considering the macroscopic na-

²²This also follows from the Bub-Clifton theorem discussed below. The sublattice of determinate quantum propositions defined by the identity and the EPR state is maximal: adding any proposition involves a contradiction.

ture of the instrument, and the rôle of the environment). (See Dickson, this vol., ch. 4, and [Bub, 1997] for further discussion.) I shall refer to this problem — the Schrödinger incompleteness of the theory — as Schrödinger’s problem. It is a *problem about truth* (or the instantiation of properties), as opposed to a distinct *problem about probabilities*.

Before formulating the probability problem, consider what was involved in the transition from classical to quantum mechanics. Quantum mechanics first appeared as Heisenberg’s matrix mechanics in 1925, following the ‘old quantum theory’, a patchwork of *ad hoc* modifications of classical mechanics to accommodate Planck’s quantum postulate. Essentially, Heisenberg modified the kinematics of classical mechanics by replacing certain classical dynamical variables, like position and momentum, with mathematical representatives — matrices — which do not commute. Shortly afterwards, Schrödinger developed a wave mechanical version of quantum mechanics and proved the formal equivalence of the two theories. It is common to understand the significance of the transition from classical to quantum mechanics in terms of ‘wave-particle duality’, the idea that a quantum system like an electron, unlike a classical system like a stone, manifests itself as a wave under certain circumstances and as a particle under other circumstances. This picture obscures far more than it illuminates. We can see more clearly what is going on conceptually if we consider the implications of Heisenberg’s move for the way we think about objects and their properties in the most general sense.

Heisenberg replaced the commutative algebra of dynamical variables of classical mechanics — position, momentum, angular momentum, energy, etc. — with a noncommutative algebra. Some of these dynamical variables take the values 0 and 1 only and correspond to properties. For example, we can represent the property of a particle being in a certain region of space by a dynamical variable that takes the value 1 when the particle is in the region and 0 otherwise. A dynamical variable like position corresponds to a set of such 2-valued dynamical variables or physical properties. In the case of the position of a particle, these are the properties associated with the particle being in region R , for all regions R . If, for all regions R , you know whether or not the particle is in that region, you know the position of the particle, and conversely. The 2-valued dynamical variables or properties of a classical system form a Boolean algebra, a subalgebra of the commutative algebra of dynamical variables.

Replacing the commutative algebra of dynamical variables with a noncommutative algebra is equivalent to replacing the Boolean algebra of 2-valued dynamical variables or properties with a non-Boolean algebra. The really essential thing about the classical mode of representation of physical systems in relation to quantum mechanics is that the properties of classical systems are represented as having the structure of a Boolean algebra or Boolean lattice. Every Boolean lattice is isomorphic to a lattice of subsets of a set.²³ To say that the properties of a classical

²³A lattice is a partially ordered set in which every pair of elements has a greatest lower bound (or infimum) and least upper bound (or supremum) with respect to the ordering, a minimum element (denoted by 0), and a maximum element (denoted by 1). A Boolean lattice is

system form a Boolean lattice is to say that they can be represented as the subsets of a set, the phase space or state space of classical mechanics. To say that a physical system has a certain property is to associate the system with a certain set in a representation space where the elements of the space — the points of the set — represent all possible states of the system. A state picks out a collection of sets, the sets to which the point representing the state belongs, as the properties of the system in that state. The dynamics of classical mechanics is described in terms of a law of motion describing how the state moves in the state space. As the state changes with time, the set of properties selected by the state changes. (For an elaboration, see [Hughes, 1995] and [Bub, 1997].)

So the transition from classical to quantum mechanics involves replacing the representation of properties as a Boolean lattice, i.e., as the subsets of a set, with the representation of properties as a certain sort of non-Boolean lattice. Dirac and von Neumann developed Schrödinger's equivalence proof into a representation theory for the properties of quantum systems as subspaces in a linear vector space over the complex numbers: Hilbert space. The non-Boolean lattice in question is the lattice of subspaces of this space. Instead of representing properties as the subsets of a set, quantum mechanics represents properties as the subspaces of a linear space — as lines, or planes, or hyperplanes, i.e., as a projective geometry. Algebraically, this is the central structural change in the transition from classical to quantum mechanics — although there is more to it: notably the fact that the state space for quantum systems is a Hilbert space over the complex numbers, not the reals, which is reflected in physical phenomena associated with the possibility of superposing states with different relative phases.

Instead of talking about properties, we can talk equivalently about propositions. (We say that a given property is instantiated if and only if the corresponding proposition is true.) In a Boolean propositional structure, there exist 2-valued homomorphisms on the structure that correspond to truth-value assignments to the propositions. In fact, each point in phase space — representing a classical state — defines a truth-value assignment to the subsets representing the propositions: each subset to which the point belongs represents a true proposition or a property that is instantiated by the system, and each subset to which the point does not belong represents a false proposition or a property that is not instantiated by the system. So a classical state corresponds to a complete assignment of truth values to the propositions, or a maximal consistent 'list' of properties of the system, and all possible states correspond to all possible maximal consistent lists.

Probabilities can be introduced on such a classical property structure as measures on the subsets representing the properties. Since each phase space point

a complemented, distributive lattice, i.e., every element has a complement (the lattice analogue of set-theoretic complementation) and the distributive law holds for the infimum and supremum. The partial ordering in a Boolean lattice represented by the subsets of a set X corresponds to the partial ordering defined by set inclusion, so the infimum corresponds to set intersection, the supremum corresponds to set union, 0 corresponds to the null set, and 1 corresponds to the set x . A Boolean algebra, defined in terms of algebraic sum (+) and product (.) operations, is equivalent to a Boolean lattice defined as a partially ordered structure.

defines a truth-value assignment, the probability of a property is the measure of the set of truth-value assignments that assign a 1 ('true') to the property — in effect, we 'count' (in the measure-theoretic sense) the relative number of state descriptions in which the property is instantiated (or the corresponding proposition is true), and this number represents the probability of the property. So it makes sense to interpret the probability of a property as a measure of our ignorance as to whether or not the property is instantiated. Probability distributions over classical states represented as phase space points are sometimes referred to as 'mixed states', in which case states corresponding to phase space points are distinguished as 'pure states.'

The problem for a quantum property structure, represented by the lattice of subspaces of a Hilbert space, arises because 2-valued homomorphisms do not exist on these structures (except in the special case of a 2-dimensional Hilbert space). If we take the subspace structure of Hilbert space seriously as the structural feature of quantum mechanics corresponding to the Boolean property structure or propositional structure of classical mechanics, the non-existence of 2-valued homomorphisms on the lattice of subspaces of a Hilbert space means that there is no partition of the totality of properties of the associated quantum system into two sets: the properties that are instantiated by the system, and the properties that are not instantiated by the system; i.e., there is no partition of the totality of propositions into true propositions and false propositions. (Of course, other ways of associating propositions with features of a Hilbert space are possible, and other ways of assigning truth values, including multi-valued truth value assignments and contextual truth value assignments. Ultimately, the issue here concerns what we take as the salient structural change involved in the transition from classical to quantum mechanics.)

It might appear that, on the standard interpretation, a pure quantum state represented by a 1-dimensional subspace in Hilbert space — a minimal element in the subspace structure — defines a truth-value assignment on quantum propositions in an analogous sense to the truth-value assignment on classical propositions defined by a pure classical state. Specifically, on the standard interpretation, a pure quantum state selects the propositions represented subspaces containing the state as true, and the propositions represented by subspaces orthogonal to the state as false. (Note that orthogonality is the analogue of set-complement, or negation, in the subspace structure; the set-theoretical complement of a subspace is not in general a subspace.)

There is, however, an important difference between the two situations. In the case of a classical state, every possible property represented by a phase space subset is selected as either instantiated by the system or not; equivalently, every proposition is either true or false. In the case of a quantum state, the properties represented by Hilbert space subspaces are not partitioned into two such mutually exclusive and collectively exhaustive sets: some propositions are assigned no truth value. Only propositions represented by subspaces that contain the state are assigned the value 'true', and only propositions represented by subspaces or-

thogonal to the state are assigned the value ‘false.’ This means that propositions represented by subspaces that are at some non-zero or non-orthogonal angle to the ray representing the quantum state are not assigned any truth value in the state, and the corresponding properties must be regarded as *indeterminate* or *indefinite*: according to the theory, there *can be no fact of the matter* about whether these properties are instantiated or not.

It turns out that there is only one way to assign (generalized) probabilities to quantum properties, i.e., weights that satisfy the usual Kolmogorov axioms for a probability measure on Boolean sublattices of the non-Boolean lattice of quantum properties. This is the content of Gleason’s theorem [Gleason, 1957]. For a quantum state ρ , a property p represented by a projection operator P is assigned the probability $\text{Tr}(\rho P)$. If ρ is a pure state $\rho = |\psi\rangle\langle\psi|$, the probability of p is $|\langle\psi_p|\psi\rangle|^2$, where $|\psi_p\rangle$ is the orthogonal projection of $|\psi\rangle$ onto the subspace P , i.e., the probability of p is the square of the cosine of the angle between the ray $|\psi\rangle$ and the subspace P . This means that properties represented by subspaces containing the state are assigned probability 1, properties represented by subspaces orthogonal to the state are assigned probability 0, and all other properties, represented by subspaces at a non-zero or non-orthogonal angle to the state are assigned a probability between 0 and 1. So quantum probabilities are not represented as measures over truth-value assignments and cannot be given an ignorance interpretation in the obvious way.

The question now is: *what do these ‘angle probabilities’ or, perhaps better, ‘angle weights’ mean?* The orthodox answer is that the probability assigned to a property of a system by a quantum state is to be understood as the probability of finding the property in a measurement process designed to ascertain whether or not that property obtains. A little thought will reveal that this proposal is very problematic. When the system is represented by a quantum state that assigns a certain property the probability 1/2, say, this property is indeterminate. Physicists would say that ascribing the property to the system in that state is ‘meaningless.’ But somehow it makes sense to design an experiment to ascertain whether or not the property is instantiated by the system. And in such a measurement, the probability is asserted to be 1/2 that the experiment will yield the answer ‘yes’, and 1/2 that the experiment will yield the answer ‘no.’ Clearly, a measurement process in quantum mechanics is not simply a procedure for ascertaining whether or not a property is instantiated in any straightforward sense. Somehow, a measurement process enables an indeterminate property, that is neither instantiated nor not instantiated by a system in a given quantum state, to either instantiate itself or not with a certain probability; or equivalently, a proposition that is neither true nor false can become true or false with a certain probability in a suitable measurement process.

The probability problem (as opposed to the truth problem, Schrödinger’s problem) is the problem of interpreting the ‘angle weights’ as probabilities in some sense (relative frequencies? propensities? subjective Bayesian betting probabilities?) that does not reduce to a purely instrumentalist interpretation of quantum

mechanics, according to which the theory is simply regarded as a remarkably accurate instrument for prediction. (Recall Einstein's remark about quantum mechanics being of interest only to shopkeepers and engineers on the Copenhagen interpretation.) The problem arises because of the unique way in which probabilities can be introduced in quantum mechanics, and because the notion of measurement or observation is utterly mysterious on the Copenhagen interpretation.

In classical theories, we measure to find out what *we don't know*, but in principle a measurement does not change *what is* (and even if it does change what is, this is simply a change or disturbance from one state of being to another that can be derived on the basis of the classical theory itself). In quantum mechanics, measurements apparently bring into being something that was *indeterminate*, not merely unknown, before, i.e., a proposition that was neither true nor false becomes true in a measurement process, and the way in which this happens according to the theory is puzzling, given our deepest assumptions about objectivity, change, and intervention.

Now, *we know how to solve Schrödinger's problem*, i.e., we know all the possible ways of modifying quantum mechanics to solve this problem. The problem arises because of the linear dynamics of the theory, which yields a certain entangled state as the outcome of a measurement interaction, and the interpretation of this entangled state as representing a state of affairs that makes certain propositions true, certain propositions false, and other propositions indeterminate. Either we change the linear dynamics in some way, or we keep the linear dynamics and say something non-orthodox about the relation between truth and indeterminateness and the quantum state. Both options have been explored in various ways and in great detail: *we understand the solution space for Schrödinger's problem, and the consequences of adopting a particular solution*.

'Collapse' theories, like the theory developed by Ghirardi, Rimini, and Weber (GRW), and extended by Pearle [Ghirardi, 2002], solve the problem by modifying the linear dynamics of quantum mechanics. (See Dickson, this vol., ch. 4, for an account.) In the modified theory, there is a certain very small probability that the wavefunction of a particle (the function defined by the components of the quantum state with respect to the position basis in Hilbert space) will spontaneously 'collapse' after being multiplied by a peaked Gaussian of a specified width. For a macroscopic system consisting of many particles, this probability can be close to 1 for very short time intervals. In effect, this collapse solution modifies the linear dynamics of standard quantum mechanics by adding uncontrollable noise. When the stochastic terms of the modified dynamics become important at the mesoscopic and macroscopic levels, they tend to localize the wave function in space. So measurement interactions involving macroscopic pieces of equipment (or cats) can be distinguished from elementary quantum processes, insofar as they lead to the almost instantaneous collapse of the wave function and the correlation of the measured observable with the position of a localized macroscopic pointer observable.

'No collapse' solutions are constrained by certain 'no go' theorems that restrict

the assignment of properties, or values to observables, under very general assumptions about the algebra of observables [Kochen and Specker, 1967], or restrict the assignment of values to observables under certain assumptions about how distributions of values are related to quantum probabilities [Bell, 1964]. A theorem by Bub and Clifton [1996] shows that if you assume that the set of definite-valued observables has a certain structure (essentially allowing quantum probabilities to be recovered as classical measures over distributions defined by different possible sets of values or properties), and the pointer observable in a measurement process belongs to the set of definite-valued observables, then the class of such theories — so-called ‘modal interpretations’ — is uniquely specified. More precisely, the sublattice associated with any single observable R is a Boolean lattice, \mathcal{B} , and a quantum state $|\psi\rangle$ defines a classical probability measure on \mathcal{B} , in the sense that all the single and joint probabilities assigned by $|\psi\rangle$ to elements in \mathcal{B} can be recovered as measures on a Kolmogorov probability space defined on the ‘phase space’ X of 2-valued homomorphisms on \mathcal{B} . The Bub-Clifton theorem characterizes the maximal lattice extension, \mathcal{L} , of any such Boolean sublattice associated with an observable R and a given quantum state $|\psi\rangle$, under the assumption that \mathcal{L} is an ortholattice,²⁴ invariant under lattice automorphisms that preserve R and $|\psi\rangle$, for which the probabilities assigned by $|\psi\rangle$ to elements in \mathcal{L} can be similarly recovered as measures on a Kolmogorov probability space defined on the ‘phase space’ Y of 2-valued homomorphisms on \mathcal{L} . In this sense, the theorem characterizes the limits of classicality in a quantum propositional structure. It turns out that different modal interpretations can be associated with different ‘determinate sublattices’ \mathcal{L} , i.e., with different choices of a ‘preferred observable’ R . For standard quantum mechanics, R is the identity, and the determinate sublattice \mathcal{L} consists of all quantum propositions represented by subspaces containing the state $|\psi\rangle$ (propositions assigned probability 1 by $|\psi\rangle$) and subspaces orthogonal to $|\psi\rangle$ (propositions assigned probability 0 by $|\psi\rangle$). Bohm’s hidden variable theory can be regarded as a modal interpretation in which the preferred observable is position in configuration space. (See Dickson, this vol., ch. 4, and [Goldstein, 2001] for an account of Bohm’s theory.)

An alternative type of ‘no-collapse’ solution to the Schrödinger problem is provided by the Everett interpretation [Everett, 1957]. (See Dickson, this vol., ch. 4, for an account.) There are a variety of Everettian interpretations in the literature, the common theme being that all possible outcomes of a measurement are regarded as actual in some indexical sense, relative to different terms in the global entangled state (with respect to a certain preferred basis in Hilbert space), which are understood to be associated with different worlds or different minds, depending on the version. The most sophisticated formulation of Everett’s interpretation is probably the Saunders-Wallace version [Saunders, 1998; Wallace, 2003]. Here the preferred basis is selected by decoherence (see below), and probabilities are introduced as rational betting probabilities in the Bayesian sense via a decision-theoretic argument originally due to Deutsch [1999].

²⁴I.e., an orthogonal complement exists for every element of \mathcal{L} .

To sum up: any solution to Schrödinger's measurement problem involves either modifying the linear dynamics of the theory ('collapse' theories), or taking some observable in addition to the identity as having a determinate value in every quantum state, and modifying what the standard theory says about what propositions are true, false, and indeterminate in a quantum state (modal interpretations, 'no collapse' hidden variable theories), so that at the end of a measurement interaction that correlates macroscopic pointer positions with possible values of a measured observable, the pointer propositions and propositions referring to measured values end up having determinate truth values. Alternatively (Everettian interpretations), we can interpret quantum mechanics so that every measurement outcome becomes determinate in some indexical sense (with respect to different worlds, or different minds, or different branches of the entangled state, etc.).

We know in considerable detail what these solutions look like, in terms of how quantum mechanics is modified. It was a useful project to explore these solutions, because we learnt something about quantum mechanics in the process, and perhaps there is more to learn by exploring the solution space further. But the point to note here is that all these solutions to the 'truth problem' of measurement distort quantum mechanics in various ways by introducing additional structural features that obscure rather than illuminate our understanding of the phenomena involved in information-theoretic applications of entanglement, such as quantum teleportation, the possibility and impossibility of certain quantum cryptographic protocols relative to classical protocols, the exponential speed-up of quantum computation algorithms relative to classical algorithms, and so on.

Consider again the Bohr-Einstein dispute about the interpretation of quantum mechanics. One might say that what separated Einstein (and Schrödinger) and Bohr was their very different answers to what van Fraassen [1991, 4] has called 'the foundational question *par excellence*: *how could the world possibly be the way quantum theory says it is?*' This would be misleading. Einstein answered this question by arguing that the world couldn't be the way quantum theory says it is, unless the theory is not the whole story (so a 'completion' of the theory — perhaps Einstein's sought-after unified field theory — would presumably answer the question). But Bohr's complementarity interpretation is not intended to be an answer to this question. Rather, complementarity should be understood as suggesting an answer to a different question: *why must the world be the way quantum theory says it is?*

To bring out the difference between these two questions, consider Einstein's distinction between what he called 'principle' versus 'constructive' theories. Einstein introduced this distinction in an article on the significance of the special and general theories of relativity that he wrote for the London *Times*, which appeared in the issue of November 28, 1919 [1919]:

We can distinguish various kinds of theories in physics. Most of them are constructive. They attempt to build up a picture of the more complex phenomena out of the material of a relatively simple formal scheme from which they start out. Thus the kinetic theory of gases seeks to

reduce mechanical, thermal, and diffusional processes to movements of molecules — i.e., to build them up out of the hypothesis of molecular motion. When we say that we have succeeded in understanding a group of natural processes, we invariably mean that a constructive theory has been found which covers the processes in question.

Along with this most important class of theories there exists a second, which I will call ‘principle theories.’ These employ the analytic, not the synthetic, method. The elements which form their basis and starting-point are not hypothetically constructed but empirically discovered ones, general characteristics of natural processes, principles that give rise to mathematically formulated criteria which the separate processes or the theoretical representations of them have to satisfy. Thus the science of thermodynamics seeks by analytical means to deduce necessary conditions, which separate events have to satisfy, from the universally experienced fact that perpetual motion is impossible.

Einstein’s point was that relativity theory is to be understood as a principle theory. He returns to this theme in his ‘Autobiographical Notes’ [1949, 51–52], where he remarks that he first tried to find a constructive theory that would account for the known properties of matter and radiation, but eventually became convinced that the solution to the problem was to be found in a principle theory that reconciled the constancy of the velocity of light in vacuo for all inertial frames of reference, and the equivalence of inertial frames for all physical laws (mechanical as well as electromagnetic):

Reflections of this type made it clear to me as long ago as shortly after 1900, i.e., shortly after Planck’s trailblazing work, that neither mechanics nor electrodynamics could (except in limiting cases) claim exact validity. By and by I despaired of the possibility of discovering the true laws by means of constructive efforts based on known facts. The longer and the more despairingly I tried, the more I came to the conviction that only the discovery of a universal formal principle could lead us to assured results. The example I saw before me was thermodynamics. The general principle was there given in the theorem: the laws of nature are such that it is impossible to construct a *perpetuum mobile* (of the first and second kind). How, then, could such a universal principle be found?

A little later [1949, 57], he adds:

The universal principle of the special theory of relativity is contained in the postulate: The laws of physics are invariant with respect to the Lorentz-transformations (for the transition from one inertial system to any other arbitrarily chosen system of inertia). This is a restricting principle for natural laws, comparable to the restricting principle

for the non-existence of the *perpetuum mobile* which underlies thermodynamics.

According to Einstein, two very different sorts of theories should be distinguished in physics. One sort involves the reduction of a domain of relatively complex phenomena to the properties of simpler elements, as in the kinetic theory, which reduces the mechanical and thermal behavior of gases to the motion of molecules, the elementary building blocks of the constructive theory. The other sort of theory is formulated in terms of ‘no go’ principles that impose constraints on physical processes or events, as in thermodynamics (‘no perpetual motion machines’). For an illuminating account of the role played by this distinction in Einstein’s work, see the discussion by Martin Klein in [1967].

The special theory of relativity is a principle theory, formulated in terms of two principles: the equivalence of inertial frames for all physical laws (the laws of electromagnetic phenomena as well as the laws of mechanics), and the constancy of the velocity of light in vacuo for all inertial frames. These principles are irreconcilable in the geometry of Newtonian space-time, where inertial frames are related by Galilean transformations. The required revision yields Minkowski geometry, where inertial frames are related by Lorentz transformations. Einstein characterizes the special principle of relativity, that the laws of physics are invariant with respect to Lorentz transformations from one inertial system to another, as ‘a restricting principle for natural laws, comparable to the restricting principle for the non-existence of the *perpetuum mobile* which underlies thermodynamics.’ (In the case of the general theory of relativity, the group of allowable transformations includes all differentiable transformations of the space-time manifold onto itself.) By contrast, the Lorentz theory [Lorentz, 1909], which derives the Lorentz transformation from the electromagnetic properties of the aether, and assumptions about the transmission of molecular forces through the aether, is a constructive theory.

The question:

How could the world possibly be the way the quantum theory says it is?

is motivated by a difficulty in interpreting quantum mechanics as a constructive theory, and the appropriate response is some constructive repair to the theory that resolves the difficulty, or the demonstration that the puzzling features of quantum mechanics at the phenomenal level (the phenomena of interference and entanglement) can be derived from a physically unproblematic constructive theory.

The question:

Why must the world be the way the quantum theory says it is?

does not ask for a ‘bottom-up’ explanation of quantum phenomena in terms of a physical ontology and dynamical laws. Rather, the question concerns a ‘top-down’ derivation of quantum mechanics as a principle theory, in terms of operational

constraints on the possibilities of manipulating phenomena. In the case of quantum mechanics, the relevant phenomena concern information.

This shift in perspective between the two questions is highlighted in a remark by Andrew Steane in his review article on ‘Quantum Computing’ [1998, 119]:

Historically, much of fundamental physics has been concerned with discovering the fundamental particles of nature and the equations which describe their motions and interactions. It now appears that a different programme may be equally important: to discover the ways that nature allows, and prevents, *information* to be expressed and manipulated, rather than particles to move.

Steane concludes his review with the following proposal [1998, 171]:

To conclude with, I would like to propose a more wide-ranging theoretical task: to arrive at a set of principles like energy and momentum conservation, but which apply to information, and from which much of quantum mechanics could be derived. Two tests of such ideas would be whether the EPR-Bell correlations thus became transparent, and whether they rendered obvious the proper use of terms such as ‘measurement’ and ‘knowledge.’

A similar shift in perspective is implicit in Wheeler’s question ‘Why the quantum?’, one of Wheeler’s ‘Really Big Questions’ [1998]. Steane’s suggestion is to answer the question by showing how quantum mechanics can be derived from information-theoretic principles. A more specific proposal along these lines originates with Gilles Brassard and Chris Fuchs. As remarked in §5.2, Brassard and Fuchs [Brassard, 2000; Fuchs, 1997; Fuchs, 2000; Fuchs and Jacobs, 2002] speculated that quantum mechanics could be derived from information-theoretic constraints formulated in terms of certain primitive cryptographic protocols: specifically, the possibility of unconditionally secure key distribution, and the impossibility of unconditionally secure bit commitment.

The CBH theorem (motivated by the Brassard-Fuchs conjecture) shows that quantum mechanics can be regarded as a principle theory in Einstein’s sense, where the principles are information-theoretic constraints. So we have an answer to the question: why must the world be the way quantum mechanics says it is? The phenomena of interference and nonlocal entanglement are bound to occur in a world in which there are certain constraints on the acquisition, communication, and processing of information.

Consider, for comparison, relativity theory, the other pillar of modern physics. A relativistic theory is a theory with certain symmetry or invariance properties, defined in terms of a group of space-time transformations. Following Einstein’s formulation of special relativity as a principle theory, we understand this invariance to be a consequence of the fact that we live in a world in which natural processes are subject to certain constraints: roughly (as Hermann Bondi [1980] puts it), ‘no overtaking of light by light’, and ‘velocity doesn’t matter’ (for electromagnetic as

well as mechanical phenomena). Recall Einstein's characterization of the special principle of relativity as 'a restricting principle for natural laws, comparable to the restricting principle of the non-existence of the *perpetuum mobile* which underlies thermodynamics.'). Without Einstein's analysis, the transformations of Minkowski space-time would simply be a rather puzzling algorithm for relativistic kinematics and the Lorentz transformation, which is incompatible with the kinematics of Newtonian space-time. What Einstein's analysis provides is a rationale for taking the structure of space and time as Minkowskian: we see that this is required for the consistency of the two principles of special relativity.

A quantum theory is a theory in which the observables and states have a certain characteristic algebraic structure. Unlike relativity theory, quantum mechanics was born as a recipe or algorithm for calculating the expectation values of observables measured by macroscopic measuring instruments. A theory with a commutative C^* -algebra has a phase space representation — not necessarily the phase space of classical mechanics, but a theory in which the observables of the C^* -algebra are replaced by 'beables' (Bell's term, see [1987]), and the C^* -algebraic states are replaced by beable-states representing complete lists of properties (idempotent quantities). In this case, it is possible to extend the theory to include the measuring instruments that are the source of the C^* -algebraic statistics, so that they are no longer 'black boxes' but constructed out of systems that are characterized by properties and states of the phase space theory. That is, the C^* -algebraic theory can be replaced by a 'detached observer' theory of the physical processes underlying the phenomena, to use Pauli's term [Born, 1971, 218], including the processes involved in the functioning of measuring instruments. Note that this depends on a representation theorem. In the noncommutative case, we are guaranteed only the existence of a Hilbert space representation of the C^* -algebra, and it is an open question whether a 'detached observer' description of the phenomena is possible.

Solving Schrödinger's problem — the truth problem — amounts to a proposal to treat quantum mechanics as a failed or incomplete constructive theory in need of constructive repair. In effect, the problem is *how to account for quantum information* — the puzzling features of interference and nonlocal entanglement — in a theoretical framework in which only classical information is meaningful in a fundamental sense. If we treat quantum mechanics as a principle theory of information, the core foundational problem is the probability problem. From this perspective, the problem is *how to account for the appearance of classical information* in a quantum world characterized by information-theoretic constraints.

One might complain that treating quantum mechanics as a principle theory amounts to simply postulating what is ultimately *explained* by a constructive theory like the GRW theory or Bohm's theory. This would amount to rejecting the idea that a principle theory can be explanatory. From the perspective adopted here, Bohm's constructive theory in relation to quantum mechanics is like Lorentz's constructive theory of the electron in relation to special relativity. Cushing [1998, 204] quotes Lorentz (from the conclusion of the 1916 edition of *The Theory of*

Electrons) as complaining similarly that ‘Einstein simply postulates what we have deduced.’

I cannot speak here of the many highly interesting applications which Einstein has made of this principle [of relativity]. His results concerning electromagnetic and optical phenomena ... agree in the main with those which we have obtained in the preceding pages, the chief difference being that Einstein simply postulates what we have deduced, with some difficulty and not altogether satisfactorily, from the fundamental equations of the electromagnetic field. By doing so, he may certainly take credit for making us see in the negative result of experiments like those of Michelson, Rayleigh and Brace, not a fortuitous compensation of opposing effects, but the manifestation of a general and fundamental principle.

Yet, I think, something may also be claimed in favour of the form in which I have presented the theory. I cannot but regard the aether, which can be the seat of an electromagnetic field with its energy and its vibrations, as endowed with a certain degree of substantiality, however different it may be from all ordinary matter. In this line of thought, it seems natural not to assume at starting that it can never make any difference whether a body moves through the aether or not, and to measure distances and lengths of time by means of rods and clocks having a fixed position relative to the aether.

Note that Lorentz’s theory is constrained by the principles of special relativity, which means that the aether as a rest frame for electromagnetic phenomena must, in principle, be undetectable. So such a theory can have no excess empirical content over special relativity. Cushing [1998, 193] also quotes Maxwell as asking whether ‘it is not more philosophical to admit the existence of a medium which we cannot at present perceive, than to assert that a body can act at a place where it is not.’ Yes, but not if we also have to admit that, in principle, as a matter of physical law, if we live in a world in which events are constrained by the two relativistic principles, the medium must remain undetectable.

You can, if you like, tell a constructive story about quantum phenomena, but such an account, if constrained by the information-theoretic principles, will have no excess empirical content over quantum mechanics. Putting this differently, a solution to Schrödinger’s truth problem that has excess empirical content over quantum mechanics must violate one or more of the CBH information-theoretic constraints. So, e.g., a Bohmian theory of quantum phenomena is like an aether theory for electromagnetic fields. Just as the aether theory attempts to make sense of the behaviour of fields by proposing an aether that is a sort of *sui generis* mechanical system different from all other mechanical systems, so Bohm’s theory attempts to make sense of quantum phenomena by introducing a field (the quantum potential or guiding field) that is a sort of *sui generis* field different from other physical fields.

The crucial distinction here is between a constructive theory formulated in terms of a physical ontology and dynamical laws (‘bottom-up’) and a principle theory formulated in terms of operational constraints at the phenomenal level (‘top-down’). A constructive theory introduces an algebra of beables and beable-states. A principle theory introduces an algebra of observables and observable-states, which are essentially probability measures.

It seems clear that the algebra of observables will be non-trivially distinct from the algebra of beables if cloning is impossible. For if a constructive theory for a certain domain of phenomena allows dynamical interactions in which a beable of

one system, designated as the measuring instrument, can become correlated with a beable of another system, designated as the measured system, without disturbing the values of other beables of the measured system, we can take such an interaction as identifying the value of the beable in question (in the sense that the value of a beable of one system is recorded in the value of a beable of a second system). If this is possible, then it will be possible to simultaneously measure any number of beables of a system by concatenating measurement interactions, and so it will be possible in principle to identify any arbitrary beable state. If we assume that we can prepare any state, then the possibility of identifying an arbitrary state means that we can construct a device that could copy any arbitrary state. So if we *cannot* construct such a device, then measurement in this sense must also be impossible. It follows that a ‘measurement’ in the constructive theory will be something other than the mere identification of a beable value of a system, without disturbance, and the question of what the observables are in such a theory will require a non-trivial analysis.

Such an analysis is indeed given by Bohm in Part II of his two-part 1952 paper on hidden variables [1952], and a more careful and sophisticated analysis is given by [Dürr *et al.*, 2003] for their ‘Bohmian mechanics’ version of Bohm’s theory. As one would expect (given the equilibrium distribution assumption which ensures that Bohm’s theory is empirically indistinguishable from quantum mechanics), while the *beables* are functions of position in configuration space (and form a commutative algebra), the *observables* of the theory are just the observables of quantum theory and form a noncommutative algebra.

The CBH theorem assumes that, for the theories we are concerned with, the *observables* form a C^* -algebra. The content of the CBH theorem is that, given certain information-theoretic constraints, the C^* -algebra of *observables* and *observable states* takes a certain form characteristic of quantum theories. The theorem says nothing about beables and beable-states, and does not address the measurement problem (Schrödinger’s truth problem), let alone solve it. But from the perspective adopted here, the measurement problem is simply the observation that cloning is impossible, and a ‘solution to the measurement problem’ is the proposal of a physical ontology and dynamics and an analysis of measurement that yields the observables and observable-states of standard quantum mechanics. Such theories provide possible explanations for the impossibility of cloning. But since there are now a variety of such explanations available, and — assuming the CBH information-theoretic principles — there are no empirical constraints, in principle, that could distinguish these explanations, there seems little point in pursuing the question further. A constructive theory whose *sole* motivation is to ‘solve the measurement problem’ seems unlikely to survive fundamental advances in physics driven by other theoretical or experimental problems

The probability problem — the core foundational problem for the interpretation of quantum mechanics as a principle theory of information — can be put this way: From the information-theoretic constraints, we get a noncommutative (or non-Boolean) theory of correlations for which there is no phase space representa-

tion. One can define, in a unique way (according to Gleason's theorem) generalized 'transition probabilities' or 'transition weights' associated with certain structural features of the noncommutative structure: the angles between geometrical elements representing quantum 'propositions.' The problem is how to understand these weights as representing probabilities, *without reducing the problem to a solution of the truth problem.*

It seems clear that we need to take account of the phenomenon of decoherence (see Landsmann, this vol., ch. 5; Dickson, this vol., ch. 4; [Zurek, 2003; Olliver *et al.*, 2004]): an extremely fast process that occurs in the spontaneous interaction between a macrosystem and its environment that leads to the virtually instantaneous suppression of quantum interference. What happens, roughly, is that a macrosystem like Schrödinger's cat typically becomes correlated with the environment — an enormous number of stray dust particles, air molecules, photons, background radiation, etc. — in an entangled state that takes a certain form with respect to a preferred set of basis states, which remain stable as the interaction develops and includes more and more particles. It is as if the environment is 'monitoring' the system via a measurement of properties associated with the preferred states, in such a way that information about these properties is stored redundantly in the environment. This stability, or robustness, of the preferred basis, and the redundancy of the information in the environment, allows one to identify certain emergent structures in the overall pattern of correlations — such as macroscopic pointers and cats and information-gatherers in general — as classical-like: the correlational information required to reveal quantum interference for these structures is effectively lost in the environment. So it appears that the information-theoretic constraints are consistent with both (i) the conditions for the existence of measuring instruments as sources of classical information, and (ii) the existence of information-gatherers with the ability to use measuring instruments to apply and test quantum mechanics, given a characterization of part of the overall system as the environment. That is, decoherence provides an explanation for the emergence of classical information in a quantum correlational structure.

If something like the above account of decoherence is acceptable, then the probability problem reduces to showing that the probabilities assigned to measurement outcomes by these information-gatherers, in the subjective Bayesian sense, are just the Gleason generalized transition probabilities. That is, we need to show that, while quantum theory, at the fundamental level, is a noncommutative theory of correlations for which there is no phase space representation, it is *also* a theory of the probabilistic behavior of information-gatherers, certain emergent structures in the pattern of correlations when correlational information in their environment is ignored. For an argument along these lines, see [Pitowsky, 2002].

On the view proposed here, no measurement outcomes are certified as determinate by the theory. Rather, measuring instruments are sources of classical information, where the individual occurrence of a particular distinguishable event ('symbol') produced stochastically by the information source lies outside the theory. In this sense, a measuring instrument, insofar as it functions as a classical

information source, is still ultimately a ‘black box’ in the theory. So a quantum description will have to introduce a ‘cut’ between what we take to be the ultimate measuring instrument in a given measurement process and the quantum phenomenon revealed by the instrument. But this ‘cut’ is no longer *ad hoc*, or mysterious, or in some other way problematic, as it is in the Copenhagen interpretation (see Landsmann, this vol., ch. 5). For here the ‘cut’ just reflects the fundamental interpretative claim: that quantum mechanics is *a theory about the representation and manipulation of information* constrained by the possibilities and impossibilities of information-transfer in our world, rather than a theory about the ways in which nonclassical waves and particles move.

ACKNOWLEDGEMENTS

This chapter was written during the tenure of a University of Maryland General Research Board semester award in 2005. I thank the University of Maryland for research support and Jeremy Butterfield for extensive and very helpful comments on a preliminary version of the chapter.

BIBLIOGRAPHY

- [Aharonov *et al.*, 1964] Y. Aharonov, P. G. Bergmann, and J. L. Lebowitz. Time-symmetric quantum states. *Physical Review B*, 134:1410–1416, 1964. Reprinted in *Quantum Theory and Measurement*, J. Archibald Wheeler and W. H. Zurek (eds.), Princeton University Press, Princeton, 680–686, 1983.
- [Aharonov *et al.*, 2005] D. Aharonov, A. Ta-Shma, U. V. Vazirani, and A. C. Yao. Quantum bit escrow. In *STOC 2000, Proceedings of the Thirty Second Annual ACM Symposium on Theory Of Computing*. *arXiv e-print quant-ph/0004017*, 2005.
- [Aspect *et al.*, 1981] A. Aspect, P. Grangier, and G. Roger. Experimental tests of realistic local theories via Bell’s theorem. *Physical Review Letters*, 47:460–477, 1981.
- [Aspect *et al.*, 1982] A. Aspect, P. Grangier, and G. Roger. Experimental realization of EPR gedankenexperiment: a new violation of Bell’s inequalities. *Physical Review Letters*, 49:91–94, 1982.
- [Bacciagaluppi, 1994] G. Bacciagaluppi. Separation theorems and Bell inequalities in algebraic quantum mechanics. In P. Busch, P. Lahti, and P. Mittelstaedt (eds.), *Symposium on the Foundations of Modern Physics 1993: Quantum Measurement, Irreversibility and the Physics of Information*, 29–37. World Scientific, Singapore, 1994.
- [Barenco, 1998] A. Barenco. Quantum computation: an introduction. In Lo *et al.* [1998], 143–183.
- [Barnum *et al.*, 1996a] H. Barnum, C. M. Caves, C. A. Fuchs, Richard Jozsa, and B. Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76:2828, 1996.
- [Barnum *et al.*, 1996b] H. Barnum, C. A. Fuchs, R. Jozsa, and B. Schumacher. General fidelity limit for quantum channels. *Physical Review A*, 54:4707, 1996.
- [Barnum *et al.*, 2001] H. Barnum, P. Hayden, R. Jozsa, and A. Winter. On the reversible extraction of classical information from a quantum source. *Proceedings of the Royal Society (London) A*, 457:2019–2039, 2001.
- [Bell, 1964] J. S. Bell. On the Einstein–Podolsky–Rosen Paradox. *Physics*, 1:195–200, 1964. Reprinted in John Stuart Bell, *Speakable and Unsayable in Quantum Mechanics*, Cambridge University Press, Cambridge, 1989.
- [Bell, 1987] J. S. Bell. Beables for quantum field theory. In *Speakable and Unsayable in Quantum Mechanics*, 173–180. Cambridge University Press, Cambridge, 1987.

- [Benioff, 1980] P. Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics*, 22(5):563–591, 1980.
- [Bennett and Brassard, 1984] C. H. Bennett and G. Brassard. Quantum cryptography: public key distribution and coin tossing. In *Proceedings of IEEE international conference on computers, systems, and signal processing*, 175–179, New York, 1984. IEEE.
- [Bennett and Wiesner, 1992] C. H. Bennett and S. J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69:2881–2884, 1992.
- [Bennett et al., 1982] C. H. Bennett, G. Brassard, S. Breidbart, and S. J. Wiesner. Quantum cryptography, or unforgeable subway tokens. In D. Chaum, R. L. Rivest, and A. T. Sherman (eds.), *Advances in Cryptology: Proceedings of Crypto 82*, 267–275, New York, 1982. Plenum Press.
- [Bennett et al., 1993] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. Wootters. Teleporting an unknown quantum state via dual classical and EPR channels. *Physical Review Letters*, 70:1895–1899, 1993.
- [Bennett, 1973] C. H. Bennett. Logical reversibility of computations. *IBM Journal of Research and Development*, 17:525–532, 1973.
- [Bohm, 1951] D. Bohm. *Quantum Theory*. Prentice Hall, Englewood Cliffs, NJ, 1951.
- [Bohm, 1952] D. Bohm. A suggested interpretation of quantum theory in terms of ‘hidden’ variables. i and ii. *Physical Review*, 85:166–193, 1952.
- [Bohr, 1949] N. Bohr. Discussion with Einstein on epistemological problems in modern physics. In P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist*, volume VII. The Library of Living Philosophers, Open Court, La Salle, IL, 1949.
- [Bondi, 1980] H. Bondi. *Relativity and Common Sense*. Dover Publications, 1980.
- [Born, 1971] M. Born. *The Born-Einstein Letters*. Walker and Co., London, 1971.
- [Boschi et al., 1998] D. Boschi, S. Branca, F. De Martini, L. Hardy, and S. Popescu. Experimental realization of teleporting an unknown pure quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 80:1121–1125, 1998.
- [Bouwmeester et al., 1997] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eible, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390:575–579, 1997.
- [Brassard et al., 1993] G. Brassard, C. Crépeau, R. Jozsa, and D. Langlois. A quantum bit commitment scheme provably unbreakable by both parties. In *Proceedings of the 34th Annual IEEE Symposium and the Foundations of Computer Science*, 362–371, New York, November 1993. IEEE.
- [Brassard, 2000] G. Brassard. Remarks on quantum foundations in the light of quantum cryptography. ‘Quantum Foundations in the Light of Quantum Information and Cryptography’, Université de Montréal, May 17–19 (2000)., 2000.
- [Brukner and Zeilinger, 2000] Č. Brukner and A. Zeilinger. Quantum measurement and Shannon information, a reply to M. J. W. Hall. *arXiv e-print quant-ph/0008091*, 2000.
- [Brukner and Zeilinger, 20002] Č. Brukner and A. Zeilinger. Information and fundamental elements of the structure of quantum theory. In *Festschrift for C. F. v. Weizsäcker on the occasion of his 90th birthday*. *arXiv e-print quant-ph/0212084*. Unknown, 20002.
- [Brukner and Zeilinger, 2001] Č. Brukner and A. Zeilinger. Conceptual inadequacy of the Shannon information in quantum measurement. *Physical Review A*, 63:022113, 2001.
- [Bub and Clifton, 1996] J. Bub and R. Clifton. A uniqueness theorem for ‘no collapse’ interpretations of quantum mechanics. *Studies in the History and Philosophy of Modern Physics*, 27:181–219, 1996.
- [Bub, 1997] J. Bub. *Interpreting the Quantum World*. Cambridge University Press, Cambridge, 1997.
- [Bub, 2001a] J. Bub. The bit commitment theorem. *Foundations of Physics*, 3:735–756, 2001.
- [Bub, 2001b] J. Bub. Secure key distribution via pre- and post-selected quantum states. *Physical Review A*, 63:032309–032311, 2001.
- [Bub, 2004] J. Bub. Why the quantum? *Studies in History and Philosophy of Modern Physics*, 35B:241–266, 2004.
- [Bub, 2005] J. Bub. Quantum mechanics is about quantum information. *Foundations of Physics*, 34:541–560, 2005.

- [Cleve *et al.*, 1998] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society A*, 454:339–354, 1998.
- [Clifton *et al.*, 2003] R. Clifton, J. Bub, and H. Halvorson. Characterizing quantum theory in terms of information-theoretic constraints. *Foundations of Physics*, 33:1561–1591, 2003.
- [Connes, 1994] A. Connes. *Noncommutative Geometry*. Academic Press, San Diego, 1994.
- [Cover and Thomas, 1991] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.
- [Cushing, 1998] J. T. Cushing. *Philosophical Concepts in Physics*. Cambridge University Press, Cambridge, 1998.
- [Deutsch and Jozsa, 1992] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London A*, 439:553, 1992.
- [Deutsch, 1985] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum compute. *Proceedings of the Royal Society of London A*, 400:497, 1985.
- [Deutsch, 1989] D. Deutsch. Quantum computational networks. *Proceedings of the Royal Society of London A*, 400:73, 1989.
- [Deutsch, 1999] D. Deutsch. Quantum theory of probability and decisions. *Proceedings of the Royal Society of London A*, 455:3129–3197, 1999.
- [Dieks, 1982] D. Dieks. Communication by EPR devices. *Physics Letters A*, 92:271–272, 1982.
- [Dürr *et al.*, 2003] D. Dürr, S. Goldstein, and N. Zanghì. Quantum equilibrium and the role of operators as observables in quantum theory. *arXiv e-print quant-ph/0308038*, 2003.
- [Einstein *et al.*, 1935] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, 1935.
- [Einstein, 1919] A. Einstein. What is the theory of relativity? *The London Times*, page 13, 1919. First published November 28, 1919. Also in A. Einstein, *Ideas and Opinions*, Bonanza Books, New York, 1954, pp. 227–232.
- [Einstein, 1949] A. Einstein. Autobiographical notes. In P. Allen Schilpp (ed.), *Albert Einstein: Philosopher-Scientist*, 3–94. Open Court, La Salle, IL, 1949.
- [Einstein, 1967] A. Einstein. Letter to Schrödinger dated december 22, 1950. In K. Przibram (ed.), *Letters on Wave Mechanics*. Philosophical Library, New York, 1967.
- [Ekert, 1991] A. Ekert. Quantum cryptography based on bell’s theorem. *Physical Review Letters*, 67:661, 1991.
- [Everett, 1957] H. Everett. ‘relative state’ formulation of quantum mechanics. *Reviews of Modern Physics*, 29:454–462, 1957.
- [Feynman, 1982] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21:467, 1982.
- [Fuchs and Jacobs, 2002] C. A. Fuchs and K. Jacobs. An information-theoretic tradeoff relation for finite-strength quantum measurements. *Physical Review A*, 63:062305–062320, 2002.
- [Fuchs, 1997] C. A. Fuchs. Information gain vs. state disturbance in quantum theory. *Fortschr. Phys.*, 46:535–565, 1997. Reprinted in *Quantum Computation: Where Do We Want to Go Tomorrow?*, S. L. Braunstein (ed.), 229–259, Wiley VCH, Weinheim (1998).
- [Fuchs, 2000] C. A. Fuchs. Just two nonorthogonal quantum states. In *Quantum Communication, Computing, and Measurement 2*, 11–16, Dordrecht, 2000. Kluwer.
- [Fuchs, 2001a] C. A. Fuchs. Notes on a Paulian idea: foundational, historical, anecdotal and forward-looking thoughts on the quantum. *arXiv e-print quant-ph/0105039*, 2001.
- [Fuchs, 2001b] C. A. Fuchs. Quantum foundations in the light of quantum information. In A. Gonis (ed.), *Proceedings of the NATO Advanced Research Workshop on Decoherence and its Implications in Quantum Computation and Information Transfer*. *arXiv e-print quant-ph/0106166*, 2001.
- [Fuchs, 2002a] C. A. Fuchs. The anti-Våxjö interpretation of quantum mechanics. *arXiv e-print quant-ph/0204146*, 2002.
- [Fuchs, 2002b] C. A. Fuchs. Quantum mechanics as quantum information (and only a little more). *arXiv e-print quant-ph/0205039*, 2002.
- [Furasawa *et al.*, 1998] A. Furasawa, J. L. Sorensen, S. L. Braunstein, C. A. Fuchs, H. J. Kimble, and E. S. Polzik. Unconditional quantum teleportation. *Science*, 282:706–709, 1998.
- [Furry, 1936] W. H. Furry. A note on the quantum mechanical theory of measurement. *Physical Review*, 49:393–399, 1936.
- [Ghirardi, 2002] G.-C. Ghirardi. Collapse theories. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, 2002. <http://plato.stanford.edu/entries/qm-collapse/>.

- [Gisin, 1989] N. Gisin. Stochastic quantum dynamics and relativity. *Helvetica Physica Acta*, 62:363–371, 1989.
- [Gleason, 1957] A. N. Gleason. Measures on the closed sub-spaces of hilbert spaces. *Journal of Mathematics and Mechanics*, 6:885–893, 1957.
- [Goldstein, 2001] S. Goldstein. Bohmian mechanics. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, 2001. <http://plato.stanford.edu/entries/qm-bohm>.
- [Grover, 1997] L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79:325, 1997.
- [Hall, 2000] M. J. W. Hall. Comment on ‘Conceptual inadequacy of the Shannon information’ by Č. Brukner and A. Zeilinger. *arXiv e-print quant-ph/0007116*, 2000.
- [Halvorson and Bub, 2005] H. Halvorson and J. Bub. Can quantum cryptography imply quantum mechanics? reply to Smolin. *Quantum Information and Computation*, 5:170–175, 2005.
- [Halvorson, 2004] H. Halvorson. Remote preparation of arbitrary ensembles and quantum bit commitment. *Journal of Mathematical Physics*, 45:4920–4931, 2004.
- [Hardy and Kent, 2004] L. Hardy and A. Kent. Cheat sensitive quantum bit commitment. *Physical Review Letters*, 92:157901, 2004. Longer version available at arXiv e-print quant-ph/9911043.
- [Hartley, 1928] R. V. L. Hartley. Transmission of information. *Bell System Technical Journal*, 7:53, 1928.
- [Howard, 2004] D. Howard. Who invented the Copenhagen interpretation? A study in mythology. <http://www.nd.edu/~dhoward1/Papers.html>, 2004.
- [Hughes, 1995] R. I. G. Hughes. *The Structure and Interpretation of Quantum Mechanics*. Cambridge University Press, 1995.
- [Hughston et al., 1993] L. P. Hughston, R. Jozsa, and W. K. Wootters. A complete classification of quantum ensembles having a given density matrix. *Physics Letters A*, 183:14–18, 1993.
- [Jaynes, 1957] E. T. Jaynes. Information theory and statistical mechanics ii. *Physical Review*, 108:171–190, 1957.
- [Jozsa and Linden, 2002] R. Jozsa and N. Linden. On the role of entanglement in quantum computational speed-up. *arXiv e-print quant-ph/0201143*, 2002.
- [Jozsa and Schumacher, 1994] R. Jozsa and B. Schumacher. A new proof of the quantum noiseless coding theorem. *Journal of Modern Optics*, 41:2343–2349, 1994.
- [Jozsa, 1997a] R. Jozsa. Entanglement and quantum computation. In S. A. Huggett, L. J. Mason, K. P. Tod, S. T. Tsou, and N. M. J. Woodhouse (eds.), *The Geometric Universe: Science, Geometry, and the Work of Roger Penrose*, 369–378. Oxford University Press, 1997.
- [Jozsa, 1997b] R. Jozsa. Quantum algorithms and the fourier transform. *arXiv e-print quant-ph/9707033*, 1997.
- [Jozsa, 1998] R. Jozsa. Quantum information and its properties. In Lo et al. [1998], 49–75.
- [Jozsa, 1999] R. Jozsa. Searching in grover’s algorithm. *arXiv e-print quant-ph/9901021*, 1999.
- [Jozsa, 2005] R. Jozsa. An introduction to measurement based quantum computation. *arXiv e-print quant-ph/0508124*, 2005.
- [Kadison and Ringrose, 1997] R. Kadison and J. Ringrose. *Fundamentals of the Theory of Operator Algebras*. American Mathematical Society, Providence, Rhode Island, 1997.
- [Kent, 1999a] A. Kent. Secure classical bit commitment using fixed capacity communication channels. *arXiv e-print quant-ph/9906103*, 1999.
- [Kent, 1999b] A. Kent. Unconditionally secure bit commitment. *Physical Review Letters*, 83:1447–1450, 1999.
- [Klein, 1967] M. J. Klein. Thermodynamics in Einstein’s thought. *Science*, 157:509–516, 1967.
- [Kochen and Specker, 1967] S. Kochen and E. P. Specker. On the problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17:59–87, 1967.
- [Landau, 1987] L. J. Landau. On the violation of Bell’s inequality in quantum theory. *Physics Letters A*, 120:54–56, 1987.
- [Landsman, 1998] N. Landsman. *Mathematical Topics Between Classical and Quantum Mechanics*. Springer, New York, 1998.
- [Lewis and Papadimitriou, 1981] H. R. Lewis and C. H. Papadimitriou. *Elements of the theory of computation*. Prentice-Hall, Englewood Cliffs, New Jersey, 1981.
- [Lo and Chau, 1997] H.-K. Lo and H. F. Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78:3410–3413, 1997.
- [Lo and Chau, 1998] H.-K. Lo and H. F. Chau. Why quantum bit commitment and ideal coin tossing are impossible. *Physica D*, 120:177–187, 1998.

- [Lo *et al.*, 1998] H.-K. Lo, S. Popescu, and T. Spiller (eds.) *Introduction to Quantum Computation and Information*, Singapore, 1998. World Scientific.
- [Lorentz, 1909] H. A. Lorentz. *The Theory of Electrons*. Columbia University Press, New York, 1909.
- [Mackey, 1963] G. Mackey. *Mathematical Foundations of Quantum Mechanics*. Benjamin, New York, 1963.
- [Mayers, 1996a] D. Mayers. The trouble with quantum bit commitment. *arXiv e-print quant-ph/9603015*, 1996.
- [Mayers, 1996b] D. Mayers. Unconditionally secure quantum bit commitment is impossible. In *Proceedings of the Fourth Workshop on Physics and Computation*, 224–228, Boston, 1996. New England Complex System Institute.
- [Mayers, 1997] D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, 1997.
- [Metzger, 2000] S. Metzger. Spin-measurement retrodiction revisited. *arXiv e-print quant-ph/0006115*, 2000.
- [Nielsen and Chuang, 2000] M. A. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [Nielsen *et al.*, 1998] M. A. Nielsen, E. Knill, and R. Laflamme. Complete quantum teleportation using nuclear magnetic resonance. *Nature*, 396:52–55, 1998.
- [Nielsen, 2003] M. A. Nielsen. Journal club notes on cluster-state quantum computation. <http://qinfo.org/qc-by-measurement/>, 2003.
- [Nielsen, 2005] M. A. Nielsen. Cluster-state quantum computation. *arXiv e-print quant-ph/0504097*, 2005.
- [Nyquist, 1924] H. Nyquist. Certain factors affecting telegraph speed. *Bell System Technical Journal*, 3:324, 1924.
- [Olliver *et al.*, 2004] H. Olliver, D. Poulin, and W. Zurek. Environment as a witness: selective proliferation of information and emergence of objectivity in a quantum universe. *Physical Review Letters*, 93:22041, 2004.
- [Pitowsky, 2002] I. Pitowsky. Betting on the outcomes of measurements: a Bayesian theory of quantum probability. *arXiv e-print quant-ph/0208121*, 2002.
- [Popescu and Rohrlich, 1998] S. Popescu and D. Rohrlich. The joy of entanglement. In Lo *et al.* [1998], 29–48.
- [Preskill, 2005] J. Preskill. Lecture notes for quantum computation. <http://www.theory.caltech.edu/people/preskill/ph229/>, 2005.
- [Raussendorf and Briegel, 2001a] R. Raussendorf and H. J. Briegel. A one-way quantum computer. *arXiv e-print quant-ph/0208121*, 2001.
- [Raussendorf and Briegel, 2001b] R. Raussendorf and H. J. Briegel. Quantum computing via measurements only. *Physical Review Letters*, 86:5188–5199, 2001.
- [Rivest *et al.*, 1978] R. L. Rivest, A. Shamir, and L. M. Adleman. A method of obtaining digital signatures and public-key cryptosystems. *Communications of the Association for Computing Machinery*, 21:120–126, 1978.
- [Saunders, 1998] S. Saunders. Time, quantum mechanics, and probability. *Synthese*, 114:373–404, 1998.
- [Schilpp, 1949] P. A. Schilpp (ed.). *Albert Einstein: Philosopher-Scientist*, volume VII, La Salle, IL, 1949. The Library of Living Philosophers, Open Court.
- [Schrödinger, 1935] E. Schrödinger. Discussion of probability relations between separated systems. *Proceedings of the Cambridge Philosophical Society*, 31:555–563, 1935.
- [Schrödinger, 1936] E. Schrödinger. Probability relations between separated systems. *Proceedings of the Cambridge Philosophical Society*, 32:446–452, 1936.
- [Schumacher, 1995] B. Schumacher. Quantum coding. *Physical Review A*, 51:2738–2747, 1995.
- [Schumacher, 1998] B. Schumacher. Lecture notes on quantum information. University of Innsbruck, 1998.
- [Shannon, 1948] C. E. Shannon. A mathematical theory of computation. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- [Shor, 1994] P. W. Shor. Algorithms for quantum computation; discrete logarithms and factoring. In *Proceedings, 35th Annual Symposium on Foundations of Computer Science*, Los Alamitos, CA, 1994. IEEE Press.
- [Shor, 1997] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal of Computation*, 26:1484–1509, 1997.

- [Simon, 1994] D. R. Simon. On the power of quantum computation. In *Proceedings, 35th Annual Symposium on Foundations of Computer Science*, 116–123, Los Alamitos, CA, 1994. IEEE Press.
- [Simon, 1997] D. R. Simon. On the power of quantum computation. *SIAM Journal of Computation*, 26:1474–1483, 1997.
- [Steane, 1998] A. Steane. Quantum computing. *Reports on Progress in Physics*, 61:117–173, 1998.
- [Steane, 2003] A. Steane. A quantum computer only needs one universe. *Studies in History and Philosophy of Modern Physics*, 34B:469–478, 2003.
- [Summers, 1990] S. J. Summers. On the independence of local algebras in quantum field theory. *Reviews in Mathematical Physics*, 2:201–247, 1990.
- [Timpson, 2004] C. Timpson. *Quantum Information Theory and the Foundations of Quantum Mechanics*. PhD thesis, Oxford, 2004.
- [Vaidman *et al.*, 1987] L. Vaidman, Y. Aharonov, and D. Z. Albert. How to ascertain the values of s_x, s_y, s_z of a spin-1/2 particle. *Phys. Rev. Letters*, 58:1385, 1987.
- [van Fraassen, 1991] B. van Fraassen. *Quantum Mechanics: An Empiricist View*. Clarendon Press, Oxford, 1991.
- [Vidal, 2003] G. Vidal. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters*, 91:147902, 2003.
- [von Neumann, 1955] J. von Neumann. *Mathematical foundations of quantum mechanics*. Princeton University Press, Princeton, 1955.
- [Wald, 1984] R. M. Wald. *General Relativity*. University of Chicago Press, 1984.
- [Wallace, 2003] D. Wallace. Everettian rationality: defending Deutsch’s approach to probability in the Everett interpretation. *Studies in the History and Philosophy of Modern Physics*, 34:415–439, 2003.
- [Wheeler, 1998] J. A. Wheeler. *Geons, Black Holes, and Quantum Foam: A Life in Physics*. W.W. Norton, New York, 1998. With K. Ford.
- [Wiesner, 1983] S. J. Wiesner. Conjugate coding. *SIGACT News*, 15:17, 1983.
- [Wooters and Zurek, 1982] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.
- [Yuen, 2005] H. P. Yuen. Unconditionally secure quantum bit commitment. *arXiv e-print quant-ph/0505132*, 2005.
- [Zurek, 2003] W. H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Review of Modern Physics*, 75:715, 2003.

THE CONCEPTUAL BASIS OF QUANTUM FIELD THEORY

Gerard 't Hooft

1 INTRODUCTION TO THE NOTION OF QUANTIZED FIELDS

Quantum Field Theory is one of those cherished scientific achievements that have become considerably more successful than they should have, if one takes into consideration the apparently shaky logic on which it is based. With awesome accuracy, all known subatomic particles appear to obey the rules of one example of a quantum field theory that goes under the uninspiring name of “The Standard Model”. The creators of this model had hardly anticipated such a success, and one can rightfully ask to what it can be attributed.

We have long been aware of the fact that, in spite of its successes, the Standard Model cannot be *exactly* right. Most quantum field theories are not asymptotically free, which means that they cannot be extended to arbitrarily small distance scales. We could try to cure the Standard Model, but this would not improve our understanding at all, because we know that, at those extremely tiny distance scales where the problems would become relevant, a force appears that we cannot yet describe unambiguously: the gravitational force. It would have to be understood first.

Perhaps this is the real strength of Quantum Field Theory: we know where its limits are, and these limits are far away. The gravitational force acting between two subatomic particles is tremendously weak. As long as we disregard *that*, the theory is perfect. And, as I will explain, its internal logic is not shaky at all.

Subatomic particles all live in the domain of physics where spins and actions are comparable to Planck’s constant \hbar . One obviously needs Quantum Mechanics to describe them. Since the energies available in sub-atomic interactions are comparable to, and often larger than, the rest mass energy mc^2 of these particles, they often travel with velocities close to that of light, c , and so relativistic effects will also be important. Thus, in the first half of the twentieth century, the question was asked:

“How should one reconcile Quantum Mechanics with Einstein’s theory of Special Relativity?”

As we shall explain, Quantum Field Theory is the answer to this question.

Our first intuitions would be, and indeed were, quite different [Pais, 1986; Crease and Mann, 1986]. One would set up abstract Hilbert spaces of states, each containing fixed or variable numbers of particles. Subsequently, one would postulate a consistent scheme of interactions. What would ‘consistent’ mean? In Quantum Mechanics, we have learned how to describe a process where we start with a certain number of particles that are all far apart but moving towards one another. This is the ‘in’ state $|\psi\rangle_{\text{in}}$. After the interaction has taken place, we end up with particles all moving away from one another, a state $|\psi'\rangle_{\text{out}}$. The probability that a certain in-state evolves into a given out-state is described by a quantum mechanical transition amplitude, ${}_{\text{out}}\langle\psi'|\psi\rangle_{\text{in}}$. The set of all such amplitudes in the vector spaces formed by all in- and out-states is called the *scattering matrix*. One can ask how to construct the scattering matrix in such a way that (i) it is invariant under Lorentz transformations, and (ii) obeys the strict laws of *quantum causality*. By ‘quantum causality’ we mean that under no circumstance a measurable effect may proceed with a velocity faster than that of light. In practice, this means that one must demand that any set of local operators $\mathcal{O}_i(\mathbf{x}, t)$ obeys commutation rules such that the commutators $[\mathcal{O}_i(\mathbf{x}, t), \mathcal{O}_j(\mathbf{x}', t')]$ vanish as soon as the vector $(\mathbf{x} - \mathbf{x}', t - t')$ is space-like. One can show that this implies that the scattering matrix must obey *dispersion relations*.

This is indeed how physicists started to think about their problem. But how should one construct such a scattering matrix? Does any systematic procedure exist?

A quantized *field* may seem to be something altogether different, yet it does appear to allow for the construction of an interacting medium that does obey the laws of Lorentz invariance and causality. The local operators can be constructed from the fields. All we then have to do is to set up schemes of relativistically covariant field equations, such as Maxwell’s laws. Even the introduction of non-linear terms in these equations appears to be straightforward, and if we were to subject such systems to a mathematically well-defined procedure called “quantization”, we would have candidates for a solution to the aforementioned problem.

Realizing that the energy in a quantized field comes in quantized energy packages, which in all respects behave like elementary particles, and, conversely, realizing that operators in the form of fields could be defined also when one starts up with Hilbert spaces consisting of elementary particles, it was discovered that quantized fields do indeed describe subatomic particles. Subsequently, it was discovered that, in a quantized field, the number of ways in which interactions can be introduced (basically by adding non-linear terms in the field equations), is quite limited. Quantization requires that all interactions can be given in the form of a Lagrange function \mathcal{L} ; relativity requires this \mathcal{L} to be Lorentz-invariant, and, most strikingly, self-consistency of Quantum Field Theory then provides further restrictions, which leads to the possibility of writing down a complete list of all possible interactions. The Standard Model is just one element of this list.

The scope of this concise treatise on Quantum Field Theory is too limited to admit detailed descriptions of all technical details. Instead, special emphasis is put on the conceptual issues that arise when addressing the numerous questions and problems associated with this doctrine. One *could* use this text to *learn* Quantum Field Theory, but for many technical details, more literature must be consulted [de Wit and Smith, 1986; Aitchison and Hey, 1989; Ryder, 1985; Itzykson and Zuber, 1980; Cheng and Li, 1984] We also limited ourselves to applications of Quantum Field Theory in elementary particle physics. There are many examples in low-temperature physics where these and similar techniques are useful, but they will not be addressed here.

2 SCALAR FIELDS

2.1 Classical Theory: Feynman rules

A field is here taken to mean a physical variable that is a function of space-time coordinates $x = (\mathbf{x}, t)$. In order for our theories to be in accordance with special relativity, we will have to specify how a field transforms under a homogeneous Lorentz transformation,

$$(1) \quad x' = Lx .$$

If a field ϕ transforms as

$$(2) \quad \phi'(x) = \phi(x') ,$$

then ϕ is called a scalar field. The *improper* Lorentz transformations, such as parity reflection P and time reversal T , are of lesser importance since we know that Nature is not exactly invariant under those.

Let us first restrict ourselves to *real* scalar fields; generalization to the case where fields are denoted by complex numbers will be straightforward. Upon quantization, scalar fields will come in energy packets that behave as spinless Bose-Einstein particles, such as π^0 , π^\pm and η^0 . Conceptually, the scalar field is the easiest to work with, but in section 9 we shall find reasons why other kinds of fields can actually improve the internal consistency of our theories.

Lorentz-invariant field equations typically take the form¹

$$(3) \quad (\partial_\mu^2 - m_{(i)}^2)\phi_i = F_i(\phi); \quad \partial_\mu^2 \equiv \vec{\partial}_x^2 - \partial_t^2 .$$

¹We use summation convention: repeated indices that are not put between brackets are automatically summed over. Greek indices μ are Lorentz indices taking 4 values, Latin indices i, j, \dots run from 1 to 3. Our metric convention is $g_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$.

Here, the index i labels different possible species of scalar fields, and $F_i(\phi)$ could be any function of the field(s) $\phi_j(x)$. Usually, however, we assume that there is a potential function $V^{\text{int}}(\phi)$, such that $F_i(\phi)$ is the gradient of V^{int} , and furthermore we assume that V^{int} is a polynomial whose degree is at most four:

$$\begin{aligned}
 V^{\text{int}}(\phi) &= \frac{1}{6}g_{ijk}\phi_i\phi_j\phi_k + \frac{1}{24}\lambda_{ijkl}\phi_i\phi_j\phi_k\phi_l; \\
 (4) \quad F_i(\phi) &= \frac{\partial V^{\text{int}}(\phi)}{\partial\phi_i} = \frac{1}{2}g_{ijk}\phi_j\phi_k + \frac{1}{6}\lambda_{ijkl}\phi_j\phi_k\phi_l,
 \end{aligned}$$

where g and λ must be totally symmetric under all permutations of their indices.¹ This is actually a limitation on the forms that $F_i(\phi)$ can take. Without this limitation, we would not have a conserved energy, and quantization of the theory would not be possible. Later, we will see why higher terms in the polynomial are not permitted (section 7).

In order to understand the general structure of the classical solutions to this set of equations, we temporarily add a function $-J_i(x)$ to $F_i(\phi)$ in Eq. (3). Subsequently, we expand the solution in powers of $J_i(x)$:

$$\begin{aligned}
 (m_{(i)}^2 - \partial_\mu^2)\phi_i(x) &= J_i(x) - \frac{\partial}{\partial\phi_i}V^{\text{int}}(\phi(x)); \\
 \phi_i(x) &= \phi_i^{(1)}(x) + \phi_i^{(2)}(x) + \psi_i^{(3)}(x) + \dots \\
 &= \int d^4y G_{ij}(x-y) \left(J_j(y) - F_j(\phi^{(1)}(y) + \phi^{(2)}(y) + \phi^{(3)}(y) + \dots) \right). \quad (5)
 \end{aligned}$$

The function $G_{ij}(x-y)$ is a solution to the equation

$$(6) \quad (m_i^2 - \partial_\mu^2)G_{ij}(x-y) = \delta_{ij}\delta(x-y),$$

while $\phi_i^{(2)}(x)$ is quadratic in $J_j(y)$ and $\phi^{(3)}(x)$ is cubic, and so on. Assembling terms of equal order in $J_j(y)$ we find a recursive procedure to solve the field equations (2.1). At the end of our calculation, we might set $J_i(x)$ equal to zero, or better, have J non-vanishing only in the far-away region where the particles originated, so that the J interaction is a simplified model for the machine that produced the particles in the far past. Indeed, in the quantum theory it will also turn out to be convenient to use J as a model for the particle detector at the end of the experiment.

We see that the solution to Eq. (2.1) can be written as the sum of a large number of terms. Each of these terms can be written in the form of a diagram, called a Feynman diagram. In these diagrams, we represent a space-time point as a dot, and the function $G_{ij}(x-y)$ as a line connecting x with y . The index i may be indicated at each line. A dot may either be associated with a term $J_i(y)$, or it is a three-point vertex associated with a coefficient g_{ijk} or a four-point vertex, going with a coefficient λ_{ijkl} . A typical Feynman diagram is sketched in Fig. 1.




Figure 1. Example of a Feynman diagram for classical scalar fields

Observe the general structure of these diagrams. There are factors $\frac{1}{2}$, $\frac{1}{6}$, etc., which can easily be read off from the symmetries of the diagram. By construction, there are no closed loops: the diagram is simply connected. This will be different in the quantized theory.

One important issue is to be addressed: the Green function, $G_{ij}(x - y)$ is not completely determined by the equation (6): one may add arbitrary combinations of the solutions of the homogeneous equation

$$(7) \quad (m_i^2 - \partial_\mu^2)G_{ij}(x - y) = 0.$$

In Fourier space, this ambiguity is reflected in the fact that one has some freedom in choosing the integration curve C in the solution²

$$(8) \quad G_{ij}(x - y) = (2\pi)^{-4} \int_C d^4k e^{ik \cdot (x-y)} \frac{\delta_{ij}}{k^2 + m_i^2}.$$

Our choice can be indicated by shifting the pole by an infinitesimal imaginary number, after which we choose the contour C to be along the real axis of all integrands. A reasonable choice is

$$(9) \quad G_{ij}^+(x - y) = (2\pi)^{-4} \int d^4k e^{ik \cdot (x-y)} \frac{\delta_{ij}}{\mathbf{k}^2 - (k^0 + i\varepsilon)^2 + m_{(i)}^2},$$

where ε is an infinitesimal, positive number. With this choice, the integration contour in the complex k^0 plane can be shifted such that the imaginary part of k^0 can be given an arbitrarily large positive value, and from this one deduces that

²An inner product $k \cdot x$ stands for $\vec{k} \cdot \vec{x} - k^0 x^0$.

the Green function will vanish as soon as the time difference, $x^0 - y^0$, is negative. This Green function, called the *forward* Green function, gives our expressions the desired causality structure: There are obviously no effects that propagate backwards in time, or indeed faster than light.

The converse choice, $G^-(x-y)$, gives us the backward solution. However, in the quantized theory, we will often be interested in yet another choice, the *Feynman propagator*, defined as

$$(10) \quad G_{ij}^F(x-y) = (2\pi)^{-4} \int d^4k e^{ik \cdot (x-y)} \frac{\delta_{ij}}{\mathbf{k}^2 - k^0{}^2 + m_i^2 - i\varepsilon},$$

where, again, the infinitesimal number $\varepsilon > 0$.

The rules to obtain the complete expansion of the solution can now be summarized as follows:

- 1) Each term can be depicted as a diagram consisting of points (vertices) connected by lines (called propagators). One end-point, \odot —, corresponds to a point x where we want to know the field ϕ ; the other end points, — \times , refer to factors $J(y^{(i)})$ for the corresponding points $y^{(i)}$, see Fig. 1.
- 2) There are no “closed loops”. *i.e.* the diagrams must be simply connected (this will be different in the quantum theory).
- 3) There are vertices with three prongs (3-vertices), $\begin{array}{c} i \\ \diagdown \\ \diagup \\ k \end{array}$, each being associated with a factor g_{ijk} , and vertices with four prongs (4-vertices), $\begin{array}{c} i \\ \diagdown \\ \diagup \\ j \\ \diagdown \\ \diagup \\ l \end{array}$, each giving a factor λ_{ijkl} .
- 4) Each line connecting two points $x^{(1)}$ and $x^{(2)}$, $x^{(1)} \bullet \xrightarrow{k} \bullet x^{(2)}$, is associated with a factor $G_{ij}(x^{(1)} - x^{(2)})$ when we work in ordinary space-time (configuration space), or a factor

$$(11) \quad \frac{\delta_{ij}}{k^2 + m_i^2 - i\varepsilon},$$

in momentum space (the reason for this $i\varepsilon$ choice will only become apparent in the quantized theory).

- 5) If we work in configuration space, we must integrate over all x values at each vertex except the one where ϕ was defined; if we work in momentum space, we must integrate over the k values, subject to the restriction of momentum conservation at each vertex: $k_{\text{out}} = \sum_{\text{in}} k_{\text{in}}$.
- 6) A ‘combinatorial factor’. For the classical theories it is $1/N$, where N is the number of permutations of the source vertices that leave the diagram unaltered.

It is not difficult to generalize the rules for the case of higher polynomials in the interactions, but this will not be needed for the time being.

2.2 Spontaneous symmetry breaking: Goldstone modes

In the classical theory, the Hamilton density is

$$(12) \quad \mathcal{H}(\mathbf{x}, t) = \frac{1}{2}\dot{\phi}_i^2 + \frac{1}{2}(\vec{\partial}\phi_i)^2 + V(\phi); \quad V(\phi) = \frac{1}{2}m_i^2\phi_i^2 + V^{\text{int}}(\phi).$$

The theory is invariant under the group of transformations

$$(13) \quad \phi'_i(x) = A_{ij}\phi_j(x),$$

if A is orthogonal and the potential function $V(\phi)$ is invariant under that group. The simplest example is the transformation $\phi \leftrightarrow -\phi$:

$$(14) \quad A = \pm 1; \quad V = V(\phi^2) = \frac{1}{2}a\phi^2 + \frac{\lambda}{24}\phi^4.$$

There are two cases to consider:

i) $a > 0$. In this case, $\phi = 0$ is the absolute minimum of V . We write

$$(15) \quad a = m^2,$$

and find that m indeed describes the mass of the particle. All Feynman diagrams have an even number of external lines. Since, in the quantum theory, these lines will be associated with particles, we find that states with an odd number of particles can never evolve into states with an even number of particles, and *vice versa*. If we define the quantum number $P_C = (-1)^N$, where N is the number of ϕ particles, then we find that P_C is conserved during interactions.

ii) $a < 0$. In this case, we see that:

- trying to identify the mass of the particle using Eq. (15) yields the strange result that the mass would be purely imaginary. Such objects (“tachyons”) are not known to exist and probably difficult to reconcile with causality, and furthermore:
- the configuration $\phi = 0$ does *not* correspond to the lowest energy configuration of the system. The lowest energy is achieved when

$$(16) \quad \phi = \pm F; \quad F^2 = -6a/\lambda.$$

It is now convenient to rewrite the potential V as

$$(17) \quad V = \frac{\lambda}{24}(\phi^2 - F^2)^2 - C,$$

where we did not bother to write down the value of the constant C , since it does not occur in the evolution equations (2.1). There are now *two equivalent vacuum states*, the minima of V . Choosing one of them, we introduce a new field variable $\tilde{\phi}$ to write

$$(18) \quad \begin{aligned} \phi &\equiv F + \tilde{\phi}; \\ V &= \frac{\lambda}{24}\tilde{\phi}^2(2F + \tilde{\phi})^2 = \frac{\lambda F^2}{6}\tilde{\phi}^2 + \frac{\lambda F}{6}\tilde{\phi}^3 + \frac{\lambda}{24}\tilde{\phi}^4, \end{aligned}$$

and we see that

- a) for the new field $\tilde{\phi}$, the mass-squared $\tilde{m}^2 = \lambda F^2/3$ is positive, and
- b) a three-prong vertex appeared, with associated factor λF . The quantum number P_C is no longer apparently conserved.

This phenomenon is called ‘spontaneous symmetry breaking’, and it plays an important role in Quantum Field Theory.

Next, let us consider the case of a continuous symmetry. The prototype example is the $U(1)$ symmetry of a complex field. The symmetry group consists of the transformations $A(\theta)$, where θ is an angle:

$$(19) \quad \Phi \equiv \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2); \quad \Phi' = A(\theta)\Phi = e^{i\theta}\Phi,$$

Again, the most general potential³ invariant under these transformations is

$$(20) \quad V(\Phi, \Phi^*) = a\Phi^*\Phi + \frac{1}{2}\lambda(\Phi^*\Phi)^2 - C,$$

In the case where the $U(1)$ symmetry is apparent, one can rewrite the Feynman rules to apply directly to the complex field Φ , noticing that one can write the potential V as a real function of the two independent variables Φ and Φ^* . With

$$(21) \quad \partial_\mu^2 \Phi = \frac{\partial V(\Phi, \Phi^*)}{\partial \Phi^*},$$

one notices that the Feynman propagators can be written with an arrow in them: an arrow points towards a point x where the function $\Phi(x)$ is called for, and away from a point x' where a factor $\Phi^*(x')$ is extracted from the potential V . At every

³Observe how we adjusted the combinatorial factors. The choices made here are the most natural ones to keep these coefficients as predictable as possible in future calculations.

vertex, as many arrows enter as they leave, and so, during an interaction, the total number of lines pointing forward in time minus the number of lines pointing backward is conserved. This is an additively conserved quantum number, to be interpreted as a ‘charge’ Q . According to *Noether’s theorem*, every symmetry is associated to such a conservation law.

However, if $a < 0$, this $U(1)$ symmetry is spontaneously broken. Then we write

$$(22) \quad V = \frac{1}{2}\lambda(\Phi^*\Phi - F^2)^2 - C, \quad F^2 \equiv -a/\lambda.$$

This time, the stable vacuum states form a closed circle in the complex plane of Φ values. Let us write

$$(23) \quad \begin{aligned} \Phi &\equiv F + \tilde{\Phi}; \quad \tilde{\Phi} \equiv \frac{1}{\sqrt{2}}(\tilde{\phi}_1 + i\tilde{\phi}_2); \\ V &= \frac{1}{2}\lambda\left(F(\tilde{\Phi} + \tilde{\Phi}^*) + \Phi^*\Phi\right)^2 \\ &= \lambda F \tilde{\phi}_1^2 + \frac{\lambda F}{\sqrt{2}}\tilde{\phi}_1(\tilde{\phi}_1^2 + \tilde{\phi}_2^2) + \frac{\lambda}{8}(\tilde{\phi}_1^2 + \tilde{\phi}_2^2)^2. \end{aligned}$$

The striking thing about this potential is that the mass term for the field $\tilde{\phi}_2$ is missing. The mass squared for the $\tilde{\phi}_1$ field is $\tilde{m}_1^2 = 2\lambda F$. The fact that one of the effective fields is massless is a fundamental consequence of the fact that we have spontaneous breakdown of a continuous symmetry. Quite generally, there is a theorem, called the *Goldstone theorem*:

If a continuous symmetry whose symmetry group has N independent generators, is broken down spontaneously into a (residual) symmetry whose group has N_1 independent generators, then $N - N_1$ massless effective fields emerge.

The propagators for massless fields obey Eq. (6) without the m^2 term, which gives these expressions an ‘infinite range’: such a Green’s function drops off only slowly for large spatial or timelike separations. These massless oscillating modes are called ‘Goldstone modes’.

2.3 Quantization of a classical theory

How does one “quantize” a field theory? In the old days of Quantum Mechanics, it was taught that “you take the Poisson brackets of the classical system, and replace these by commutators.” Here and there, one had to readjust the *order* in which classical expressions emerge, when they are replaced by operators, but the rules appeared to leave no essential ambiguities. Indeed, *if* such a procedure is possible, one may get a quantum theory which reproduces the original classical system in the limit of vanishing \hbar . Also, the group of symmetry transformations under which the classical system was invariant, often re-emerges in the quantum system.

A *field theory*, however, has a strictly infinite set of physical degrees of freedom (the field values at every point in 3-space, or, the complete set of Fourier modes). More often than not, upon “quantization”, this leads to infinities that render the theory ill-defined. One has to formulate the notion of “quantization” much more carefully, going through several intermediate steps. Since, today, the answers to our questions are so well known, it is often forgotten how these answers can be derived rigorously and why they take the form they have. What is the strictly logical sequence of arguments?

First of all, it is unreasonable to expect that every classical field theory should have a quantum mechanical counterpart. What we wish to do, is construct *some* quantum system, its Hilbert space and its Hamiltonian, such that in one or more special limits, it reproduces a known classical theory. We demand certain properties of the theory, such as Lorentz invariance and causality, but most of all we demand that it be internally logically impeccable, allowing us to calculate how in such a system particles interact, under all imaginable circumstances. We will, however, continue to use the phrase ‘quantization’, meaning that we attempt to construct a quantum theory with a given classical field theory as its $\hbar \rightarrow 0$ limit.

Often, authors forget to mention the first, very important, step in this logical procedure: replace the classical field theory one wishes to quantize by a strictly finite theory. Assuming that physical structures smaller than a certain size will not be important for our considerations, we replace the continuum of three-dimensional space by a discrete but dense lattice of points. In the differential equations, we replace all derivatives $\partial/\partial x^i$ by finite ratios of differences: $\Delta/\Delta x^i$, where $\Delta\phi$ stands for $\phi(x + \Delta x) - \phi(x)$. In Fourier space, this means that wave numbers \vec{k} are limited to a finite range (the Brillouin zone), so that integrations over \vec{k} can never diverge.

If this lattice is sufficiently dense, the solutions we are interested in will hardly depend on the details of this lattice, and so, the classical system will resume Lorentz invariance and the speed of light will be the practical limit for the velocity of perturbances. If necessary, we can also impose periodic boundary conditions in 3-space, and in that case our system is completely finite. Finite systems of this sort allow for ‘quantization’ in the old-fashioned sense: replace the Poisson brackets by commutators. Note that we did not (yet) discretize time, so the Hamiltonian of our theory has the form

$$\begin{aligned}
 H &= T + V \\
 &= \sum_{x^a} \prod_{a=1}^3 (\Delta x^a) \left(\frac{1}{2} \sum_i (\partial\phi_i/\partial t)^2 + \frac{1}{2} \sum_{i,a} \left(\frac{\Delta\phi_i}{\Delta x^a} \right)^2 + V(\phi) \right). \quad (24)
 \end{aligned}$$

The canonical momenta associated to the fields $\phi_i(x)$ are

$$(25) \quad p^i(x) = (\partial\phi_i/\partial t) \prod_{a=1}^3 (\Delta x^a),$$

and so, we will assume these to be operators obeying:

$$(26) \quad [\phi_i(x), \phi_j(x')] = 0 \quad [p^i(x), p^j(x')] = 0; \quad [\phi_i(x), p^j(x')] = i\delta_i^j \delta_{x, x'}.$$

Now, we have to wait and see what happens in the limit of an infinitely dense space-lattice. Will, like the classical theory, our quantum concoction turn out to be Lorentz-invariant? How do we perform Lorentz transformations on physical states? This question turns out to be far from trivial to answer, but the answer is known. We first need some useful technical tools.

2.4 The Feynman path integral

The Feynman path integral is often introduced as an “infinite dimensional” integral. Again, we insist on at first keeping everything finite. Label the generalized coordinates (here the ϕ_i fields) as q_i . The momenta are p_i . The Hamiltonian (2.3) is of the conventional type (the volume elements $\prod_{a=1}^3 (\Delta x^a)$ act as masses). For future use, we need a slightly more general one, a Hamiltonian that also contains pieces linear in the momenta p_i :

$$(27) \quad H = T + V; \quad T = \sum_i \frac{(p_i - A_i(\mathbf{q}))^2}{2m_{(i)}}; \quad V = V(\mathbf{q}).$$

In principle, we keep the number n of coordinates and momenta finite, in which case there is no doubt that the differential equations in question have unique, finite solutions (assuming the functions A_i and V to be sufficiently smooth; indeed we will mostly work with polynomials). Consider the configuration states $|\mathbf{q}\rangle$ and the momentum states $|\mathbf{p}\rangle$. We have

$$(28) \quad \langle \mathbf{q} | \mathbf{q}' \rangle = \delta^n(\mathbf{q} - \mathbf{q}'), \quad \langle \mathbf{p} | \mathbf{p}' \rangle = \delta^n(\mathbf{p} - \mathbf{p}'); \quad \langle \mathbf{q} | \mathbf{p} \rangle = (2\pi)^{-n/2} e^{ip_i q_i}.$$

Taking the order of the operators into account, we write for the kinetic energy

$$(29) \quad T = \sum_i \frac{p_i^2 - 2A_i p_i + A_i^2}{2m_{(i)}} + iW(\mathbf{q});$$

$$W(\mathbf{q}) = \sum_i \frac{[A_i(\mathbf{q}), p_i]}{2im_{(i)}} = \sum_i \frac{\partial_i A_i(\mathbf{q})}{2m_{(i)}}.$$

This enables us to compute swiftly the matrix elements

$$(30) \quad \langle \mathbf{q} | H | \mathbf{p} \rangle = \langle \mathbf{q} | \mathbf{p} \rangle (h(\mathbf{q}, \mathbf{p}) + iW(\mathbf{q}));$$

$$(31) \quad \langle \mathbf{p} | H | \mathbf{q} \rangle = \langle \mathbf{p} | \mathbf{q} \rangle (h(\mathbf{q}, \mathbf{p}) - iW(\mathbf{q})),$$

where $h(\mathbf{q}, \mathbf{p})$ is the *classical* Hamiltonian as a function of the two sets of variables \mathbf{q} and \mathbf{p} .

The *evolution operator* $U(t, \delta t)$ for a short time interval δt is

$$(32) \quad U(t, \delta t) = e^{-iH(t)\delta t} = \mathbf{I} - iH \delta t + \mathcal{O}(\delta t)^2.$$

Its matrix elements between states $\langle \mathbf{p} |$ and $| \mathbf{q} \rangle$ are easy to derive now:

$$\begin{aligned} \langle \mathbf{p} | U(t, \delta t) | \mathbf{q} \rangle &= \langle \mathbf{p} | \mathbf{q} \rangle - i\delta t \langle \mathbf{p} | H | \mathbf{q} \rangle + \mathcal{O}(\delta t)^2 \\ &= (2\pi)^{-n/2} e^{-i\mathbf{p} \cdot \mathbf{q}_i} \left(1 - i\delta t \{ h(\mathbf{q}, \mathbf{p}) - iW(\mathbf{q}) \} + \mathcal{O}(\delta t)^2 \right) \\ &= (2\pi)^{-n/2} \exp \left(-i\mathbf{p} \cdot \mathbf{q} - i\delta t \{ h(\mathbf{q}, \mathbf{p}) - iW(\mathbf{q}) \} + \mathcal{O}(\delta t)^2 \right). \end{aligned} \quad (33)$$

What makes this expression very useful is the fact that it does not become singular in the limit $\delta t \downarrow 0$. The momentum-momentum and the coordinate-coordinate matrix elements do become singular in that limit.

Next, let us consider a finite time interval T . The evolution operator over that time interval can formally be viewed as a sequence of many evolution operators over short time intervals δt , with $T = N \delta t$. Using closure, both in \mathbf{p} space and in \mathbf{q} space, at all time intervals,

$$(34) \quad \mathbf{I} = \int d^n \mathbf{q} | \mathbf{q} \rangle \langle \mathbf{q} | = \int d^n \mathbf{p} | \mathbf{p} \rangle \langle \mathbf{p} |,$$

we can write

$$\begin{aligned} | \psi(\mathbf{q}_N, T) \rangle &= \langle \mathbf{q}_N | U(0, T) | \psi(0) \rangle = \int d^n \mathbf{q}_0 \int d^n \mathbf{p}_0 \cdots \int d^n \mathbf{q}_{N-1} \int d^n \mathbf{p}_{N-1} \\ &\quad \langle \mathbf{q}_N | \mathbf{p}_{N-1} \rangle \langle \mathbf{p}_{N-1} | U(t_{N-1}, \delta t) | \mathbf{q}_{N-1} \rangle \langle \mathbf{q}_{N-1} | \mathbf{p}_{N-2} \rangle \\ &\quad \langle \mathbf{p}_{N-2} | U(t_{N-2}, \delta t) | \mathbf{q}_{N-2} \rangle \cdots \langle \mathbf{p}_0 | U(0, \delta t) | \mathbf{q}_0 \rangle \langle \mathbf{q}_0 | \psi(0) \rangle. \end{aligned} \quad (35)$$

Plugging in Eq. (2.4), we see that

$$\begin{aligned} | \psi(\mathbf{q}_N, T) \rangle &= \left(\prod_{\tau=0}^{N-1} \int d^n \mathbf{q}_\tau \int d^n \mathbf{p}_\tau \frac{e^{-W(\mathbf{q}_\tau)\delta t}}{(2\pi)^n} \right) \times \\ (36) \quad &\exp i \sum_{\tau=0}^{N-1} \delta t \left(\mathbf{p}_\tau \frac{\mathbf{q}_{\tau+1} - \mathbf{q}_\tau}{\delta t} - h(\mathbf{q}_\tau, \mathbf{p}_\tau, t_\tau) \right) \langle \mathbf{q}_0 | \psi(0) \rangle. \end{aligned}$$

Define

$$(37) \quad \dot{\mathbf{q}}_\tau \equiv \frac{\mathbf{q}_{\tau+1} - \mathbf{q}_\tau}{\delta t},$$

and

$$(38) \quad L(\mathbf{p}, \mathbf{q}, \dot{\mathbf{q}}, t) = \mathbf{p} \cdot \dot{\mathbf{q}} - h(\mathbf{q}, \mathbf{p}, t),$$

and the *measure*

$$(39) \quad \prod_{\tau=0}^{N-1} \int d^n \mathbf{q}_\tau \int d^n \mathbf{p}_\tau \frac{e^{-W(\mathbf{q}_\tau) \delta t}}{(2\pi)^n} \equiv \int \mathcal{D}\mathbf{q} \mathcal{D}\mathbf{p},$$

then we obtain an expression that seems to be easy to extend to infinitely fine grids in the time variable:

$$(40) \quad \langle \mathbf{q}_N | \psi(T) \rangle = \int \mathcal{D}\mathbf{q} \mathcal{D}\mathbf{p} \left(\exp i \sum_{\tau=0}^{N-1} \delta t L(\mathbf{p}, \mathbf{q}, \dot{\mathbf{q}}, t) \right) \langle \mathbf{q}_0 | \psi(0) \rangle.$$

In these expressions, we actually allowed the parameters in the Hamiltonian H and the Lagrangian L to depend explicitly on time t , so as to expose the physical structure of these expressions. Note that

$$(41) \quad L(\mathbf{p}, \mathbf{q}, \dot{\mathbf{q}}, t) = - \sum_i \frac{(p_i - A_i - m_{(i)} \dot{q}_i)^2}{2 m_{(i)}} - V(\mathbf{q}) + \sum_i (A_i \dot{q}_i + \frac{1}{2} m_{(i)} \dot{q}_i^2),$$

and the integrals over all momentum variables are easy to perform, giving some constant that only depends on the masses $m_{(i)}$:

$$(42) \quad \langle \mathbf{q}_N | \psi(T) \rangle = \int \mathcal{D}\mathbf{q} \exp \left(i \sum_{\tau=0}^{M-1} \delta t L(\mathbf{q}, \dot{\mathbf{q}}, t) \right) \langle \mathbf{q}_0 | \psi(0) \rangle,$$

with

$$(43) \quad \begin{aligned} L(\mathbf{q}, \dot{\mathbf{q}}, t) &= T - V; \quad T = \sum_i \left(\frac{1}{2} m_{(i)} \dot{q}_i^2 + A_i \dot{q}_i \right); \\ \mathcal{D}\mathbf{q} &= e^{-\sum_\tau W(\mathbf{q}_\tau) \delta t} \prod_{\tau=0}^{N-1} \left(d^n \mathbf{q}_\tau \prod_i \left(\frac{m_{(i)}}{2\pi \delta t} \right)^{\frac{1}{2}} \right). \end{aligned}$$

Actually, $L(\mathbf{q}, \dot{\mathbf{q}}, t)$ is obtained from $L(\mathbf{p}, \mathbf{q}, \dot{\mathbf{q}}, t)$ by *extremizing* the latter with respect to \mathbf{p} :

$$(44) \quad \frac{\partial}{\partial p_i} L(\mathbf{p}, \mathbf{q}, \dot{\mathbf{q}}, t) = 0; \quad \dot{q}_i = \frac{\partial h(\mathbf{q}, \mathbf{p}, t)}{\partial p_i}.$$

This is exactly the standard relation between Lagrangian and Hamiltonian of the classical theory. So, L is indeed the Lagrangian.

If the continuum limit exists, the exponent in Eq. (42) is exactly i times the classical action,

$$(45) \quad S = \int dt L(\mathbf{q}, \dot{\mathbf{q}}, t).$$

It is tempting to assume that the $\mathcal{O}(\delta t)^2$ terms in Eq. (32) disappear in the limit; after all, they are only multiplied by factors $N \approx C/\delta t$. In that case, the evolution operator in Eq. (42) clearly takes the form of an integral over all paths going from \mathbf{q}_0 to \mathbf{q}_N . This is Feynman's path integral. In the case of a field theory, one considers the field defined on a lattice in space, and since the path integral starts with a lattice in the time variable, we end up dealing with a lattice in space and time. In conclusion:

The evolution operator in a field theory is described by first rephrasing the theory on a dense lattice in space-time. Replacing partial derivatives by the corresponding finite difference ratios, one writes an expression for the action S of the theory. Normally, it can be written as an integral over a Lagrange density, $\mathcal{L}(\phi, \partial_\mu \phi)$. The evolution operator of the theory is obtained by integrating e^{iS} over all field configurations $\phi(\mathbf{x}, t)$ in a given space-time patch. The integration measure is defined from Eq. (43).

The A_i terms, linear in the time derivatives, do not play a role in scalar field theories but they do in vector theories, and the fact that they occur in the measure (43) is usually ignored. Indeed, in most cases, $W(\mathbf{q})$ vanishes, but we must be aware that it might cause problems in some special cases. We ignore the W term for the time being.

2.5 Feynman rules for the quantized theory

The Feynman rules for quantized field theories were first derived by careful analysis of perturbation theory. Writing the quantum Hamiltonian H as $H = H_0 + H^{\text{int}}$, one assembles all terms bilinear in the fields and their derivatives in H_0 and performs the perturbation expansion for small values of H^{int} . This leads to a set of calculation rules very similar to the rules derived for a classical theory, see subsection 2.1. Most of these rules (but not everything) can now most elegantly be derived from the path integral.

Let us first derive these rules for computing a finite dimensional integral of the type (42). Although often our action will not contain terms linear in the variables $q_i(t)$, we do need such terms now, so, if necessary, we add them by hand, only to

remove them at the end of the calculations. There is no need to indicate the time variable t explicitly; we absorb it in the indices i . The action is then

$$\begin{aligned}
 S(\mathbf{q}) &= \sum_{\mathbf{x}, t} \mathcal{L}(\mathbf{x}, t) \\
 &= J_i q_i - \frac{1}{2} M_{ij} q_i q_j - \frac{1}{6} A_{ijk} q_i q_j q_k - \frac{1}{24} B_{ijkl} q_i q_j q_k q_l. \tag{46}
 \end{aligned}$$

To calculate $\int d^N \mathbf{q} e^{iS(\mathbf{q})}$, we keep only the bilinear part (the term with the coefficients M_{ij}) inside the exponent, and expand the exponent of all other terms:

$$\begin{aligned}
 \text{out} \langle 0|0 \rangle_{\text{in}} &= C \int d^N \mathbf{q} \left(\exp \left(-\frac{1}{2} i M_{ij} q_i q_j \right) \right) \sum_{k=0}^{\infty} \sum_{\ell=0}^{\infty} \sum_{m=0}^{\infty} \frac{1}{k! \ell! m!} \times \\
 &\quad (i J_{i_1} q_{i_1}) \cdots (i J_{i_k} q_{i_k}) \left(-\frac{i}{6} A_{i_1 j_1 k_1} q_{i_1} q_{j_1} q_{k_1} \right) \cdots \left(-\frac{i}{6} A_{i_\ell j_\ell k_\ell} q_{i_\ell} q_{j_\ell} q_{k_\ell} \right) \\
 (47) \quad &\quad \left(-\frac{i}{24} B_{i_1 j_1 k_1 \ell_1} q_{i_1} q_{j_1} q_{k_1} q_{\ell_1} \right) \cdots \left(-\frac{i}{24} B_{i_m j_m k_m \ell_m} q_{i_m} q_{j_m} q_{k_m} q_{\ell_m} \right).
 \end{aligned}$$

(C is a constant not depending on the coefficients, but only on their dimensionality).

We can calculate all of these integrals if we know how to do the J terms. These however can be done to all orders since we know exactly how to do the Gaussian integral

$$\begin{aligned}
 \int d^N \mathbf{q} \exp i \left(-\frac{1}{2} M_{ij} q_i q_j + J_i q_i \right) &= \frac{(2\pi)^{\frac{N}{2}}}{(\det(M))^{\frac{1}{2}}} \exp \left(\frac{1}{2} i J_i M_{ij}^{-1} J_j \right) = \\
 (48) \quad C \sum_{k=0}^{\infty} \frac{1}{k!} &\left(\frac{1}{2} i J_{i_1} M_{i_1 j_1}^{-1} J_{j_1} \right) \cdots \left(\frac{1}{2} i J_{i_k} M_{i_k j_k}^{-1} J_{j_k} \right).
 \end{aligned}$$

This expression tells us how to do the integrals in Eq. (47) by collecting terms that go with given powers of J_i . The outcome of this calculation can be summarized in a concise way:

- 1) Each term can be depicted as a diagram consisting of points (vertices) connected by lines (propagators). The lines may end at points i , $\text{---} \times$, which refer to factors J_i .

- 2) There are vertices with three prongs (3-vertices), $\begin{matrix} i & & j \\ & \diagdown & / \\ & & k \end{matrix}$, each being associated with a factor A_{ijk} , and vertices with four prongs (4-vertices), $\begin{matrix} & & j & & \\ & \diagdown & / & \diagdown & / \\ i & & & & k \\ & \diagup & \diagup & \diagdown & \diagdown \\ & & l & & \end{matrix}$, each giving a factor B_{ijkl} .

- 3) Each line connecting two points i and j , is associated with a factor M_{ij}^{-1} .

- 4) In contrast with the classical theory, however, the diagrams may contain disconnected pieces, or multiply connected parts: closed loops. See Fig. 2.
- 5) There are combinatorial factors arising from the coefficients such as $k!$ in Eq. (47). One can gain experience in deriving these factors; they follow directly from the symmetry structure of a diagram. This technical detail will not be further addressed here.

Apparently nothing changes if one re-inserts the (\mathbf{x}, t) dependence of these coefficients, when the variables q_i are replaced by the fields $\phi_i(\mathbf{x}, t)$, and the action by that of a field theory:

$$S = \int d^4x \mathcal{L}(\mathbf{x}, t);$$

$$(49) \quad \mathcal{L}(\mathbf{x}, t) = -\frac{1}{2}(\partial_\mu \phi_i)^2 - \frac{1}{2}m_{(i)}^2 \phi_i^2 - V(\phi) + J_i(x)\phi_i(x).$$

The rules are as in Subsection 2.1, with the only real distinction that, in the quantum theory, diagrams with closed loops in them contribute. These diagrams may be regarded as the “quantum corrections” to the classical field theory. The disconnected diagrams mentioned under point (4), arise for technical reasons that we will not further elaborate; in practical calculations they may usually be ignored.




Figure 2. Example of a Feynman diagram for quantized scalar fields

At one point, however, we made an omission: the overall constant C was not computed. It comes from the cancellation of two coefficients (the one in the measure and the one coming from the Gaussian integrals) each of which tend to infinity in the limit of an infinitely dense grid. In most cases, we are not interested in this coefficient (it refers to vacuum-energy), but this does imply that more is needed to extract relevant physical information from these Feynman diagrams. Fortunately, this deficit is easy to cure. The “source insertions”, $J_i(x)\phi_i(x)$ can serve as a model both for the production and for the detection of particles. Let both $|0\rangle_{\text{in}}$ and $|0\rangle_{\text{out}}$ be the vacuum, or ground state of the theory. At early times, the insertion $-J(\mathbf{x}, t)\phi$ in the Hamiltonian acts on this vacuum state to excite it into the initial state we are interested in. By differentiating with respect to J ,

we can reach any initial state we want to consider. Similarly, at the end of the experiment, at late times, $J\phi$ can link the particle state that we wish to detect to the final vacuum state. In short, differentiating with respect to $J(\mathbf{x}, t)$ gives us any matrix element that we wish to study. This is easier than one might think: J_i refers to particles of type i , and if we give it the same space-time dependence as the wave function of the particle we want to see (put it on the ‘mass shell’ of that particle), then we can be sure that there will be no contamination from unwanted particle states. One only has to check the normalization, but also that is not hard: we adjust the 1-particle to 1-particle amplitude to be one; a single particle cannot scatter (it could be unstable, but that is another matter). The constant C always drops out of these calculations.

An important point is the ambiguity of the inverse matrix M^{-1} . As in the classical case, there are homogeneous solutions, so, if we work in momentum space, there will be the question how to integrate around the poles of the propagator. The $i\varepsilon$ prescription mentioned in subsection 2.1 is now imperative. This is explained as follows. Consider the propagator in position space, and choose its poles situated as follows:

$$(50) \quad \int d^4k \frac{e^{i\mathbf{k}\cdot\mathbf{x} - ik^0t}}{m^2 + \mathbf{k}^2 - k^0{}^2 - i\varepsilon}; \quad \varepsilon \downarrow 0.$$

The poles are at $k^0 = \pm(\sqrt{m^2 + \mathbf{k}^2} - i\varepsilon)$. Now consider this propagator at time $t = -T + i\beta$ with both T and β large. Since β is large, the choice of the contour at negative k^0 is immaterial, since the exponential there is very small. At positive k^0 , we choose the contour to go above the pole, so the imaginary part of k^0 is chosen positive. We see that then the exponential vanishes rapidly at negative time. In short, our propagator tends to zero if the time t tends to $-T + i\beta$ when both T and β are large and positive. The same holds for $t \rightarrow +T - i\beta$. Indeed, we want our evolution operator to be dominated by the empty diagram in these two limits. Write:

$$(51) \quad \langle \psi | U(0, +T - i\beta) | \psi' \rangle = \sum_E \langle \psi | E \rangle \exp(-iET - \beta E) \langle E | \psi' \rangle,$$

where $|E\rangle$ are the energy eigenstates. At large β , the vacuum state should dominate. Conversely, if we consider evolution *backwards* in time, the other $i\varepsilon$ prescription is needed. One then works with the Feynman rules for the inverse, or the complex conjugate, of the scattering matrix.

Now, we are in a position to add the prescription how to identify the external lines (the lines sticking out of the diagram) with in- and out-going particles. For an ingoing particle, we use a source function $J(x)$ whose Fourier components emit a *positive* amount of energy k^0 . For an out-going particle the source emits a *negative* k^0 . According to the rules formulated above, these sources would be connected to the rest of the diagram by propagators, in Fourier space $(k_\mu^2 + m^2 - i\varepsilon)^{-1}$. Since

the in- and out-going particles have $k_\mu^2 + m^2 = 0$, we must take the *residue of the pole*. In practice, this means that we have to remove the external propagators, a procedure called ‘amputation’. One then still has to establish a normalization factor. This factor is most easily obtained by checking unitarity of the scattering matrix, using the optical theorem. At first sight, this seems to be just a simple numerical coefficient, but there is a slight complication at higher orders, when self-energy corrections affect the propagator. These corrections also remove *unstable particles* from the physical scattering matrix. We return to this in Section 6. The complete Feynman rules are listed in subsection 4.5.

3 SPINOR FIELDS

3.1 The Dirac equation

The fields introduced in the previous section can only be used to describe particles with spin 0. In a quantum theory, particles can come in any representation of the *little group*, which is the subgroup of the inhomogeneous Lorentz group that leaves the 4-momentum of a particle unaffected. For massive particles in ordinary space, this is the group of rotations of a three-vector, $SO(3)$. Its representations are labelled by either an integer ≥ 0 , or an integer $+\frac{1}{2}$, representing the total spin of a particle. So, next in line are the particles with spin $\frac{1}{2}$. The wave function for such a particle has two components, one for spin up and one for spin down. Therefore, to describe a relativistic theory with such particles, we should use a two-component field obeying a relativistically covariant field equation. Paul Dirac was the first to find an appropriate relativistically covariant equation for a free particle with spin $\frac{1}{2}$:

$$(52) \quad (m + \sum_{\mu} \gamma^{\mu} \partial_{\mu}) \psi(x) = 0,$$

but the field $\psi(\mathbf{x}, t)$ has *four complex* components. Here, γ^{μ} , $\mu = 0, 1, 2, 3$, are four 4×4 matrices, obeying

$$(53) \quad \{\gamma^{\mu}, \gamma^{\nu}\} = \gamma^{\mu} \gamma^{\nu} + \gamma^{\nu} \gamma^{\mu} = 2g^{\mu\nu}; \quad \gamma_{\mu}^{\dagger} = g_{\mu\nu} \gamma^{\nu}.$$

In contrast to the scalar case, the Dirac equation is *first order* in the space- and time-derivatives, and furthermore, one could impose a ‘reality condition’ (Majorana condition) on the fields, of the form

$$(54) \quad \psi(x) = C \psi^*(x), \quad \gamma^{\mu} C = C (\gamma^{\mu})^*, \quad \mu = 0, 1, 2, 3.$$

These two features combined give the Dirac field the same multiplicity as two scalar fields. Usually, we do not impose the Majorana condition, so that the Dirac

field is truly complex, having a conserved $U(1)$ charge much like two complex scalar fields.

We briefly recapitulate the most salient features of the Dirac equation. The 4×4 Dirac matrices can conveniently be expressed in terms of two commuting sets of Pauli matrices, σ_a and τ_a . Define

$$(55) \quad \sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

and similarly for the τ matrices, except that they act in different spaces: a Dirac index is then viewed as a pair $(i\alpha)$ of indices i and α , such that the matrices σ_a act on the first index i , and the matrices τ_A act on the indices α . We have:

$$(56) \quad \sigma_a \sigma_b = \delta_{ab} + i\varepsilon_{abc} \sigma_c, \quad \tau_A \tau_B = \delta_{AB} + i\varepsilon_{ABC} \tau_C, \quad [\sigma_a, \tau_B] = 0.$$

Define (with the convention $g_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$)

$$(57) \quad \gamma^1 = \sigma_1 \tau_1, \quad \gamma^2 = \sigma_2 \tau_1, \quad \gamma^3 = \sigma_3 \tau_1, \quad \gamma^0 = -i\tau_3.$$

The matrix C in Eq. (54) is then:

$$(58) \quad C = \gamma_2 \gamma_4.$$

In the non-relativistic limit, the Dirac equation reads

$$(59) \quad (m + i\gamma^\mu k^\mu)\psi \approx (m - i\gamma^0 k^0)\psi \approx m(1 - \tau_3)\psi = 0,$$

so that only two of the four field components survive (those with $\tau_3|\psi\rangle = |\psi\rangle$). This continues to be the case for relativistic particles, simply because of Lorentz-invariance.

3.2 Fermi-Dirac statistics

At this point, we could now attempt to pursue our fundamental quantization program: produce the Poisson brackets of the system, replace these by commutators, rewrite the Hamiltonian of the system in operator form, and solve the resulting Schrödinger equation.

Unfortunately, if one uses ordinary (commuting) numbers, this does not work. The Lagrangian associated to the Dirac equation will read

$$(60) \quad L = \int d^3\vec{x} \mathcal{L}(x); \quad \mathcal{L}(x) = -\bar{\psi}(x)(m + \sum_{\mu=0}^4 \gamma^\mu \partial_\mu)\psi(x),$$

and the canonical procedure would give as momentum fields:

$$(61) \quad p_\psi(\vec{x}) = \frac{\partial \mathcal{L}}{\partial(\partial_0 \psi(\vec{x}))} = \bar{\psi}(\vec{x})\gamma^0, \quad p_{\bar{\psi}}(\vec{x}) = 0.$$

From this, one finds the Hamiltonian:

$$(62) \quad H = \int d^3\vec{x} \mathcal{H}(\vec{x}); \quad \mathcal{H}(\vec{x}) = p_\psi \dot{\psi} - \mathcal{L}(\vec{x}) = \bar{\psi}(x)(m + \sum_{i=1}^3 \gamma^i \partial_i) \psi(x).$$

Here, the index i is a spatial one, running from 1 to 3. This, however, is not bounded from below! Such a quantum theory would not possess a vacuum state, and hence be unsuitable as a model for Nature.

For a better understanding of the situation, we strip the Dirac equation to its bare bones. After diagonalizing it, we find that the Lagrangian consists of elementary units of the form

$$(63) \quad L = \bar{\psi}(i\partial_t \psi - M\psi); \quad p_\psi = i\bar{\psi}; \quad H = \bar{\psi}M\psi.$$

If we were using ordinary numbers, the only way to obtain a lower bound on H would be by identifying $\bar{\psi}$ with ψ . Then, however, the kinetic part of the Lagrangian would become a time-derivative:

$$(64) \quad \bar{\psi}\partial_t \psi \rightarrow \frac{1}{2}\partial_t(\bar{\psi}\psi),$$

so that it could not contribute to the action. One concludes that, only in the space of *anticommuting numbers*, can the Lagrangian (63) make sense. Thus, one replaces the Poisson brackets for ψ and $\bar{\psi}$ by *anticommutators*:

$$(65) \quad \{\bar{\psi}, \psi\} \equiv \bar{\psi}\psi + \psi\bar{\psi} = 1; \quad \{\psi, \psi\} = 0; \quad \{\bar{\psi}, \bar{\psi}\} = 0.$$

The elementary representation of this algebra is in a ‘Hilbert space’ consisting of just two states (the empty state and the one-particle state), in which the operators ψ and $\bar{\psi}$ act as annihilators and creators:

$$(66) \quad \psi = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}; \quad \bar{\psi} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}; \quad H = \begin{pmatrix} 0 & 0 \\ 0 & M \end{pmatrix}.$$

Returning to the non-diagonalized case, we can keep the Lagrangian (60) and Hamiltonian (62) when the commutation rules (65) are replaced by

$$(67) \quad \{\bar{\psi}^i(\mathbf{x}), \psi_j(\mathbf{x}')\} = \delta_j^i \delta(\mathbf{x} - \mathbf{x}'); \\ \{\psi_i(\mathbf{x}), \psi_j(\mathbf{x}')\} = 0; \quad \{\bar{\psi}^i(\mathbf{x}), \bar{\psi}^j(\mathbf{x}')\} = 0.$$

The commutation rules (67) turn Dirac particles into fermions. It appears to be a condition for any Lorentz-invariant quantum theory to be consistent, that *integer* spin particles must be *bosons* and particles whose spin is an integer $+\frac{1}{2}$ must be *fermions*.

3.3 The path integral for anticommuting fields

Let us now extend the notion of path integrals to include Dirac fields. This means we have to integrate over anticommuting numbers, to be called θ_i , where i is some index (possibly including \mathbf{x}). They are numbers, not operators, so all anticommutators vanish. Consider the Taylor expansion of a function of a variable θ . Since $\theta^2 = 0$, this expansion has only two coefficients:

$$(68) \quad f(\theta) = f(0) + f'(0)\theta.$$

So, this is the most general function of θ that one can have. It is generally agreed that one should define integrals for anticommuting numbers θ by postulating

$$(69) \quad \int d\theta 1 \equiv 0; \quad d\theta \theta \equiv 1.$$

The reason for this definition is that one can manipulate these expressions in the same way as integrals over ordinary numbers:

$$(70) \quad \int d\theta f(\theta + \alpha) = \int d\theta f(\theta); \quad \int d\theta \frac{\partial f(\theta)}{\partial \theta} = 0,$$

etc.

Now, consider the Hamiltonian for just one fermionic degree of freedom, (66), which we write as

$$(71) \quad H = M b^\dagger b; \quad \{b, b^\dagger\} = 1; \quad b^2 = (b^\dagger)^2 = 0,$$

and a wave function $\psi = \begin{pmatrix} \psi_0 \\ \psi_1 \end{pmatrix}$. Define the following function of θ :

$$(72) \quad \psi(\theta) \equiv \psi_0 \theta + \psi_1,$$

This now serves as our wave function. It is not hard to derive how the annihilation operator b and the creation operator b^\dagger act on these wave functions:

$$(73) \quad \text{if } \phi = b\psi \quad \text{then} \quad \phi(\theta) = \theta\psi(\theta),$$

or:

$$(74) \quad b = \theta; \quad b^\dagger = \frac{\partial}{\partial \theta}.$$

We now wish to express the evolution of a fermionic wave function in terms of a path integral, just as in subsection 2.4. Consider a short time interval δt . Then, ignoring all terms of order $(\delta t)^2$, one derives

$$\begin{aligned} e^{-i\delta t H} \psi(\theta_1) &= \psi_0 \theta_1 + (1 - iM \delta t) \psi_1 \\ &= \int d\theta_0 (-\theta_1 + \theta_0 - iM \delta t \theta_0) (\psi_0 \theta_0 + \psi_1) \\ &= \int d\theta_0 \int d\bar{\theta} \left(1 + \bar{\theta} (-\theta_1 + \theta_0 - iM \delta t \theta_0) \right) \psi(\theta_0) \\ (75) \quad &= \int d\theta_0 \int d\bar{\theta} e^{\bar{\theta}(-\theta_1 + \theta_0 - iM \delta t \theta_0)} \psi(\theta_0). \end{aligned}$$

Repeating this procedure over many infinitesimal time intervals, with $T = N \delta t$, one arrives at the formal expression

$$\begin{aligned} \psi(\theta_T) &= \int d\theta_{T-1} d\bar{\theta}_{T-1} \cdots d\theta_0 d\bar{\theta}_0 \\ &\quad \exp \sum_{\tau=0}^{N-1} \delta t \left(\bar{\theta}_\tau \left(\frac{-\theta_{\tau+1} + \theta_\tau}{\delta t} - iM \theta_\tau \right) \right) \psi(\theta_0). \end{aligned} \quad (76)$$

The exponential tends to

$$(77) \quad i \int dt L(t).$$

Thus, as in the bosonic case, the evolution operator is formally the path integral of e^{iS} over all (anticommuting) fields $\psi_i(\mathbf{x}, t)$, where the action S is the time integral of the Lagrangian L , and indeed the space-time integral of the Lagrange density $\mathcal{L}(\mathbf{x}, t)$.

In some applications, careful considerations of the boundary conditions for Dirac's equation, require an extra boundary term to be added to the action (77). In our present treatment this is of no consequence.

3.4 The Feynman rules for Dirac fields

Let M_{ij} be any matrix that can be diagonalized. Using Eqs. (69), we find the integral

$$(78) \quad \prod_i \int d\theta_i \int d\bar{\theta}_i e^{\bar{\theta}_i M_{ij} \theta_j} = \det_{ij}(M),$$

which can be easily checked by diagonalizing M , and writing

$$(79) \quad \int d\theta \int d\bar{\theta} e^{\bar{\theta}M\theta} = \int d\theta \int d\bar{\theta} (1 + \bar{\theta}M\theta) = M.$$

Thus, a Gaussian integral over anticommuting numbers gives a result very similar to that over commuting numbers, except that we get $\det(M)$ rather than $C/\det(M)$. Writing

$$\begin{aligned} M &= M_0 + \delta M; \\ \det(M) &= e^{\text{Tr}(\log M)} \\ &= 1 + \text{Tr}(\log M) + \frac{1}{2}(\text{Tr} \log M)^2 + \dots \\ (80) \quad \text{Tr} \log(M) &= \text{Tr} \log(M_0) + \text{Tr} \log(1 + M_0^{-1}\delta M), \end{aligned}$$

we see that this can be obtained from $\det(M^{-1})$ by switching the signs of all odd terms in this expansion. Since the N^{th} term corresponds to a Feynman diagram with N closed fermionic loops, one derives that the Feynman rules can be read off from the ones for ordinary commuting fields, by switching a sign whenever a closed fermionic loop is encountered.

We have

$$(81) \quad -\text{Tr} \log M = -\text{Tr} \log M_0 + \sum_{n=1}^{\infty} \frac{(-1)^n}{n} \text{Tr} (M_0^{-1}\delta M)^n.$$

Here, as in the bosonic case, $-M_0$ is the propagator of the theory, and δM represents the contribution from any perturbation. Thus, if our Lagrangian, including possible interaction terms, is

$$(82) \quad \mathcal{L} = -\bar{\psi}_i(m_{(i)} + \gamma^\mu \partial_\mu)\psi_i + \bar{\psi}_i g_{ij}(\phi)\psi_j,$$

then the propagator, in Fourier space, is

$$(83) \quad (m_{(i)} + i\gamma^\mu k_\mu)^{-1} = \frac{m_{(i)} - i\gamma^\mu k_\mu}{m_{(i)}^2 + k^2 - i\varepsilon},$$

while $g_{ij}(\phi)$ generates the interaction vertices of a Feynman diagram. The $i\varepsilon$ term is chosen as in bosonic theories, for the same reason as there: the vacuum state must be the state with lowest energy.

The poles in the propagator can be used to define in- and out-going particles, by adding source terms to the Lagrangian:

$$(84) \quad \delta\mathcal{L} = \bar{\eta}(x)\psi + \bar{\psi}\eta(x),$$

where $\eta(x)$ and $\bar{\eta}(x)$ are kept fixed, as anticommuting numbers. We could proceed to derive the precise rules for in- and out-going particles with spin up or down, but it is more convenient to postpone this until we discuss the unitarity property of the S -matrix, where these rules are required explicitly, and where we find the precise prescription for the normalization of these states (section 6).

Note that our Lagrangian is always kept to be bilinear in the anticommuting fields. This is because we insist that \mathcal{L} itself must be a commuting number and, furthermore, terms that are *quartic* in the fermionic fields have too high a dimension. We will see in the following section why such terms have to be avoided.

4 GAUGE FIELDS

4.1 Renormalizability

We continue to search for *elementary* fields, whose Lorentz covariant field equations can be subject to our quantization program. In principle, such fields could come as any arbitrary representation of the Poincaré algebra, that is, we might consider any kind of tensor field, $A_{\mu\nu\lambda\dots}(\mathbf{x}, t)$. It turns out, however, that tensors with more than one Lorentz index cannot be used. This is because we wish the *energy density* of a field to be bounded from below, and in addition, we wish the *dimensionality* of the interactions to be sufficiently low, such that all coupling strengths have mass dimension zero or positive.

A theory is called “renormalizable” if *all* of its interaction parameters λ_i (that is, all parameters with respect to which we need to make a perturbation expansion) have a mass-dimensionality that is positive or zero. In practice, the dimensionality of coupling coefficients is easy to establish; this is further explained in Section 7, *Renormalization*. Coupling strengths with mass dimension less than zero give rise to unacceptably divergent expressions for the contributions of the interactions at short scales. A prime example of a field one would *like* to include is the gravitational field described by the metric $g_{\mu\nu}(x)$, but its only possible interaction is the gravitational one, whose coupling strength, Newton’s constant G_N , has the wrong dimension. The non-renormalizable theories one then obtains are the subject of intense investigations but fall outside the scope of this paper (see C. Rovelli’s contribution in this book).

So, only spin-one fields $A_\mu^a(x)$ are left for consideration. Here, μ is a Lorentz index, while the number of field types is counted by the index $a = 1, \dots, N_V$. These fields should describe the creation and annihilation of spin-one particles. When at rest, such a particle will be in one of three possible spin states. Yet, to be Lorentz-invariant, a vector field A_μ should have four components. One of these, at least, should therefore be unphysical, although one might think of accepting an extra, spinless particle to be associated to the vector particles. More important

therefore is the consideration that, in the corresponding classical theory, the energy should be bounded from below.

This then rules out the treatment of a four-vector field as if we had four scalar fields, because the Lorentz-invariant product has an indefinite metric. Can we construct a Lagrangian for a vector field that gives a Hamiltonian that is bounded from below?

Let us look at the high-momentum limit for one of these vector fields. The only two terms in a Lagrangian that can survive there are:

$$(85) \quad \mathcal{L} = -\frac{1}{2}\alpha (\partial_\mu A_\nu)^2 + \frac{1}{2}\beta \partial_\mu A_\mu \partial_\nu A_\nu,$$

since other terms of this dimensionality can be reduced to these ones by partial integration of the action, while mass terms (terms without partial derivatives) become insignificant. We have for the canonical momentum fields

$$(86) \quad \begin{aligned} E^i &= \frac{\partial \mathcal{L}}{\partial \partial_0 A_i} = \alpha \partial_0 A_i \quad (i = 1, 2, 3); \\ E^0 &= \frac{\partial \mathcal{L}}{\partial \partial_0 A_0} = (\beta - \alpha) \partial_0 A_0 - \beta \partial_i A_i. \end{aligned}$$

Now, consider the Hamiltonian density $\mathcal{H} = E^\mu \partial_0 A_\mu - \mathcal{L}$. It must be bounded from below for all field configurations $A_\mu(\mathbf{x}, t)$. Let us first consider the case when the spacelike components A_i and all spacelike derivatives ∂_i are negligible compared to $\partial_0 A_0$:

$$(87) \quad \mathcal{H} \rightarrow \frac{1}{2}(\beta - \alpha)(\partial_0 A_0)^2,$$

then, when A_0 and all time-derivatives are negligible:

$$(88) \quad \mathcal{H} \rightarrow \frac{1}{2}\alpha(\partial_i A_j)^2 - \frac{1}{2}\beta(\partial_i A_i)^2.$$

These must all be bounded from below. Eq. (87) dictates that $\beta \geq \alpha$, while Eq. (88) dictates that $\alpha \geq \beta$. We conclude that $\alpha = \beta$, which we can both normalize to one. Since total derivatives in the Lagrangian do not count, we can then rewrite the original Lagrangian (85) as

$$(89) \quad \mathcal{L} \rightarrow -\frac{1}{4}F_{\mu\nu}^a F_{\mu\nu}^a, \quad F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a.$$

Realizing that this is the Lagrangian for ordinary QED, we know that its energy-density is properly bounded from below. We conclude that *every* vector field theory must have a Lagrangian that approaches Eq. (89) at high energies and momenta.

We do note, that with the choice $\alpha = \beta$, both (87) and (88) tend to zero. Indeed, any field A_μ^a that can be written as a space-time gradient, $A_\mu^a = \partial_\mu \Lambda^a(\mathbf{x}, t)$, has

$F_{\mu\nu}^a = 0$, and hence contributes neither to the Lagrangian nor to the Hamiltonian. Such fields could be arbitrarily strong, yet carry zero energy. They would represent particles and forces without energy. This is unacceptable in a decent Quantum Field Theory. How do we protect our theory against such features?

There is exactly one way to do this. We must make sure that field replacements of the type

$$(90) \quad A_\mu^a \rightarrow A_\mu^a + \partial_\mu \Lambda^a(x) + \dots,$$

do not affect at all the physical state that we are describing. This is what we call a local gauge transformation. We must insist that our theory is invariant under local gauge transformations. The ellipses in Eq. (90) indicate that we allow extra terms that do not contribute to the bilinear part of the Lagrangian (89). Thus, we arrive at Yang-Mills field theory.

4.2 The Yang-Mills equations

Our conclusion from the above is that every vector field is associated to a *local gauge symmetry*. The dimensionality of the local gauge group must be equal to N_V , the number of vector fields present. Besides the vector fields, the local symmetry transformations may also affect the scalar and spinor fields. In short, the vector fields must be Yang-Mills fields. We here give a brief summary of Yang-Mills theory [Yang and Mills, 1954].

We have a local Lie group with elements $\Omega(x)$ at the point x . Let the matrices T^a , $a = 1, \dots, N_V$ be its infinitesimal generators:

$$(91) \quad \Omega(x) = \mathbf{I} + i \sum_a \Lambda^a(x) T^a; \quad T^a = (T^a)^\dagger.$$

Characteristic for the group are its structure constants f_{abc} :

$$(92) \quad [T^a, T^b] = i f_{abc} T^c.$$

As is well-known in group theory, one can choose the normalization of T^a in such a way that the f_{abc} are totally antisymmetric:

$$(93) \quad f_{abc} = -f_{bac} = f_{bca}.$$

Usually, the spinor fields $\psi(x)$ and scalar fields $\phi(x)$ are introduced in such a way that they transform as (sets of irreducible) *representations* of the gauge group. A *local gauge transformation* is then:

$$(94) \quad \psi'(x) = \Omega(x)\psi(x); \quad \phi'(x) = \Omega(x)\phi(x),$$

and in infinitesimal form:

$$(95) \quad \psi'(x) = \psi(x) + i\Lambda^a(x)T^a\psi(x) + \mathcal{O}(\Lambda)^2,$$

and similarly for $\phi(x)$. The *dimension* of the irreducible representation can be different for different field types. So, scalar and spinor fields usually form gauge-vectors of various dimensionalities. In these, and in the subsequent expressions, the indices labelling the various components of the fields ψ , Ω and T^a have been suppressed.

Our vector fields $A_\mu^a(x)$ are most conveniently introduced by demanding the possibility of constructing gauge-covariant gradients of these fields:

$$(96) \quad D_\mu\psi(x) \equiv (\partial_\mu + igA_\mu^a(x)T^a)\psi(x),$$

where g is a freely adjustable coupling parameter. The repeated indices a , denoting the different species of vector fields, are to be summed over. By demanding the transformation rule

$$(97) \quad (D_\mu\psi(x))' = \Omega(x)D_\mu\psi(x) = D_\mu\psi(x) + i\Lambda^a(x)T^aD_\mu\psi(x) + \mathcal{O}(\Lambda)^2,$$

one easily derives the transformation rule for the vector fields $A_\mu^a(x)$:

$$\begin{aligned} igA_\mu^{a'}(x)T^a &= \Omega(x)\left(\partial_\mu + igA_\mu^a(x)T^a\right)\Omega^{-1}(x) \\ &= igA_\mu^a(x)T^a - i\partial_\mu\Lambda^a(x)T^a + g[T^a, T^b]\Lambda^a(x)A_\mu^b(x) \end{aligned} \quad (98)$$

(omitting the $\mathcal{O}(\Lambda)^2$ terms). With Eq. (92), this becomes

$$(99) \quad A_\mu^{a'}(x) = A_\mu^a(x) - \frac{1}{g}\partial_\mu\Lambda^a(x) + f_{abc}\Lambda^b(x)A_\mu^c(x).$$

If we ensure that all gradients used are covariant gradients, we can directly construct the general expressions for Lagrangians for scalar and spinor fields that are locally gauge-invariant:

$$(100) \quad \mathcal{L}_{\text{scalar}}^{\text{inv}}(x) = -\frac{1}{2}(D_\mu\phi)^2 - V(\phi^2);$$

$$(101) \quad \mathcal{L}_{\text{Dirac}}^{\text{inv}}(x) = -\bar{\psi}(\gamma^\mu D_\mu + m)\psi,$$

and in addition other possible invariant local interaction terms without derivatives.

The commutator of two covariant derivatives is

$$\begin{aligned} [D_\mu, D_\nu]\psi(x) &= igF_{\mu\nu}^a(x)T^a\psi(x); \\ (102) \quad F_{\mu\nu}^a(x) &= \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf_{abc}A_\mu^b A_\nu^c. \end{aligned}$$

Unlike $A_\mu^a(x)$ or the direct gradients of $A_\mu^a(x)$, this *Yang-Mills field* $F_{\mu\nu}^a$ transforms as a true adjoint representation of the local gauge group:

$$(103) F_{\mu\nu}^a{}'(x) = F_{\mu\nu}^a(x) + f_{abc}\Lambda^b(x)F_{\mu\nu}^c(x).$$

This allows us to construct a locally gauge invariant Lagrangian for the vector field:

$$(104) \mathcal{L}_{\text{YM}}^{\text{inv}}(x) = -\frac{1}{4}F_{\mu\nu}^a(x)F_{\mu\nu}^a(x).$$

The structure constants f_{abc} in the definition 102 of the field $F_{\mu\nu}^a$ implies the presence of interaction terms in the Yang-Mills Lagrangian 104. If f_{abc} is non-vanishing, we talk of a *non-Abelian* gauge theory.

There is one important complication in the case of fermions: the Dirac matrix $\gamma^5 \equiv \gamma^1\gamma^2\gamma^3\gamma^4$ can be used to project out the *chiral sectors*:

$$(105) \psi \equiv \psi_L + \psi_R; \quad \psi_L = \frac{1}{2}(1 + \gamma^5)\psi; \quad \psi_R = \frac{1}{2}(1 - \gamma^5)\psi.$$

Since the kinetic part of a Dirac Lagrangian can be split according to

$$(106) \mathcal{L}_{\text{Dirac}} = -\bar{\psi}_L(\gamma D)\psi_L - \bar{\psi}_R(\gamma D)\psi_R,$$

we may choose the left-handed fields ψ_L to be in representations different from the right-handed ones, ψ_R . However, since a mass term joins left to right:

$$(107) -m\bar{\psi}\psi = -m\bar{\psi}_L\psi_R - m\bar{\psi}_R\psi_L,$$

such terms would then be forbidden, hence such chiral fields must be massless. Secondly, not all combinations of chiral fermions are allowed. An important restriction is discussed in section 8. The fields ψ_L turn out to describe spin- $\frac{1}{2}$ massless particles with only the left-rotating helicity, while their antiparticles, described by $\bar{\psi}_L$, have only the right-rotating helicity.

4.3 The need for local gauge-invariance

In the early days of Gauge Theory, it was thought that local gauge-invariance could be an 'approximate' symmetry. Perhaps one could add mass terms for the vector field that violate local symmetry, but make the model look more like the observed situation in particle physics. We now know, however, that such models suffer from a serious defect: they are non-renormalizable. The reason is that renormalizability requires our theory to be consistent up to the very tiniest distance scales. A mass term would, at least in principle, turn the field configurations described by the $\Lambda(x)$ contributions in Eq. (99) into physically observable fields (the Lagrangian

now does depend on $\Lambda(x)$). But, since the kinetic term for $\Lambda(x)$ is lacking, violently oscillating Λ fields carry no sizeable amount of energy, so they would not be properly suppressed by energy conservation. Uncontrolled short distance oscillations are the real, physical cause for a theory being non-renormalizable.

It is similar uncontrolled short-distance fluctuations of the space-time metric that cause the quantized version of General Relativity (“Quantum Gravity”) to be non-renormalizable. Drastic measures (String Theory?) are needed to repair such a theory.

Since renormalizability provides the required coherence of our theories, local gauge symmetry, described by Eqs. (94) and (99), *must* be an exact, not an approximate symmetry of any Quantum Field Theory.⁴ Obviously, the fact that most vector particles in the sub-atomic world do carry mass must be explained in some other way. It is here that the Brout-Englert-Higgs mechanism comes to the rescue, see Section 5.

4.4 Gauge fixing

The longitudinal parts of the vector fields do not occur directly in the Yang-Mills Lagrangian (104), exactly because of its invariance under transformations of the form (99). Yet if we wish to describe solutions, we need to choose a longitudinal component. This is why we wish to impose some additional constraint, the so-called gauge condition, on our description of the solutions, both in the classical and in the quantized theory. In electrodynamics, we usually impose a constraint such as $\partial_\mu A_\mu(x) = 0$ or $A_0 = 0$. In a Yang-Mills theory, such a constraint is needed for each value of the index a . A gauge fixing term is indicated by a field $C^a(x)$ which is put equal to zero:

$$(108) \quad C^a(x) = 0; \quad a = 1, \dots, N_V; \quad \text{where}$$

$$(109) \quad \text{either } C^a(x) = \partial_\mu A_\mu^a(x) \quad (\text{Feynman gauge}),$$

$$(110) \quad \text{or } C^a(x) = A_0^a(x) \quad (\text{timelike gauge}),$$

or other possible gauge choices. It is always possible to find a $\Lambda^a(x)$ that obeys one of such conditions. For instance, to obtain the Feynman gauge (109), all one has to do is extremize an integral under variations of the gauge group:

$$(111) \quad \delta \int d^4x \left(A_\mu^a(x) \right)^2 = 0 \quad \rightarrow \quad \partial_\mu A_\mu^a(x) = 0.$$

For the classical theory, the most elegant way to impose such a gauge condition

⁴One apparent exception could be the case where the longitudinal component decouples completely, which happens in massive QED. But even in that case, it is better to view the longitudinal photon as a Higgs field, see section 5.

is by adding a Lagrange multiplier term to the Lagrangian:

$$(112) \mathcal{L}(x) = \mathcal{L}^{\text{inv}}(x) + \lambda^a(x)C^a(x),$$

where $C^a(x)$ is any of the possible gauge fixing terms and $\lambda^a(x)$ a free kinematical variable. Here, \mathcal{L}^{inv} stands for the collection of all gauge-invariant terms in the Lagrangian. The Euler-Lagrange equations of the theory then automatically yield the Yang-Mills field equations plus the constraint, apart from a minor detail: the boundary condition. Varying the gauge transformations, one finds, since \mathcal{L}^{inv} does not vary, $D_\mu \lambda^a(x) = 0$. We need to impose the stricter equation $\lambda^a(x) = 0$, which is obtained by imposing $\lambda^a(x) = 0$ at the boundaries of our system.

Alternatively, one can replace the invariant Lagrangian by

$$(113) \mathcal{L}(x) = \mathcal{L}^{\text{inv}}(x) - \frac{1}{2} \left(C^a(x) \right)^2,$$

which has the advantage that, after partial integration, the bilinear part becomes very simple: $\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F_{\mu\nu} - \frac{1}{2}(\partial_\mu A^\mu)^2 \rightarrow -\frac{1}{2}(\partial_\mu A_\nu)^2$, so that the vector field can be treated as if it were just 4 scalars. Again, varying the gauge transformation $\Lambda^a(x)$, one finds $D_\mu C^a(x) = 0$, which must be replaced by the more stringent condition $C^a(x) = 0$ by adding the appropriate boundary condition.

Note that the Lagrange-Hamilton formalism could give the wrong sign to the energy of some field components; we should continue to use the energy deduced before imposing the gauge constraint. If we use the timelike gauge (110), the energy is correct, but the theory appears to lack Lorentz invariance. Lorentz transformations must now be accompanied by gauge transformations.

How is the gauge constraint to be handled in the quantized theory? This problem was solved by B.S. DeWitt [1964; 1967a; 1967b] and by Faddeev and Popov [1967; 1969; 1984]. The gauge constraint is to be imposed in the integrand of the functional integral:

$$(114) Z = \int \mathcal{D}A(x) \int \mathcal{D}\phi(x) \dots e^{i \int d^4x \mathcal{L}^{\text{inv}}(x)} \prod_{a,x} \delta(C^a(x)) \Delta\{A, \phi\}.$$

Thus, we integrate only over those field configurations that obey the gauge condition. $\Delta\{A, \phi\}$ is a Jacobian factor, which we will discuss in a moment. The formal delta function can be replaced by a Lagrange multiplier:

$$(115) \int \mathcal{D}\lambda^a(x) e^{i \int d^4x \lambda^a(x) C^a(x)},$$

and indeed, if $\lambda^a(x)$ is simply added to the list of dynamical field variables of the theory, the Feynman rules can be derived unambiguously as they were for the scalar and the spinor case.

There is, however, a problem. It appears to be difficult to prove gauge-invariance. More precisely: we need to ascertain that, if we make the transition to a different gauge fixing function $C^a(x)$, the physical contents of the theory, in particular the scattering matrix, remains the same. The difficulty has to do with the *measure* of the integral. It is not gauge-invariant, unless we add the extra term $\Delta\{A, \phi\}$ in Eq. (114). This term is associated to the volume of an infinitesimal gauge transformation. Suppose that the field combination $C^a(x)$ transforms under a gauge transformation as

$$(116) \quad C^{a'}(x) = C^a(x) + \frac{\partial C^a(x)}{\partial \Lambda^b(x')} \Lambda^b(x'),$$

then the required volume term is the Jacobian

$$(117) \quad \Delta\{A, \phi\} = \det \left(\frac{\partial C^a(x)}{\partial \Lambda^b(x')} \right).$$

The determinant is computed elegantly by using the observation in subsection 3.4 that a Gaussian integral over anticommuting variables gives a determinant (Eq. (78)). So, we introduce anticommuting scalar fields η and $\bar{\eta}$, and then write

$$(118) \quad (117) = \int \mathcal{D}\eta^a(x) \int \mathcal{D}\bar{\eta}^a(x) \exp \left(\bar{\eta}^a(x) \frac{\partial C^a(x)}{\partial \Lambda^b(x')} \eta^b(x') \right).$$

This is called the *Faddeev-Popov* term in the action. Taking everything together, we arrive at the following action for a Yang-Mills theory:

$$(119) \quad \mathcal{L}(x) = \mathcal{L}^{\text{inv}}(x) + \lambda^a(x) C^a(x) + \bar{\eta}^a(x) \frac{\partial C^a(x)}{\partial \Lambda^b(x')} \eta^b(x').$$

It is also possible to find the quantum analogue for the classical Lagrangian (113). First, replace $C^a(x)$ by $C^a(x) - F^a(x)$, where $F^a(x)$ is a fixed but x -dependent quantity in the functional integral (119). Physical effects should be completely independent of $F^a(x)$. Therefore, we can functionally integrate over $F^a(x)$, using any weight factor we like. Choose the weight factor $e^{-\frac{1}{2} \int d^4x (F^a(x))^2}$. The Lagrange multiplier $\lambda^a(x)$ now simply forces $C^a(x)$ to be equal to $F^a(x)$. We end up with the effective Lagrangian⁵

$$(120) \quad \mathcal{L}(x) = \mathcal{L}^{\text{inv}}(x) - \frac{1}{2} \left(C^a(x) \right)^2 + \bar{\eta}^a(x) \frac{\partial C^a(x)}{\partial \Lambda^b(x')} \eta^b(x').$$

This is the most frequently used Lagrangian for gauge theories. In contrast to the Lagrangians for scalar and spinor fields, not all fields here represent physical particles. The longitudinal part of the vector fields, and the fermionic yet scalar fields $\bar{\eta}$ and η are “ghosts”.

⁵One usually absorbs the factor $1/g$ of Eq. (99) into the definition of the $\bar{\eta}$ field.

4.5 Feynman rules

The Feynman rules, needed for the computation of the scattering matrix elements using perturbation theory, can be read off directly from the gauge-fixed Lagrangian (119) or (120). In both cases, we first split off the bilinear parts⁶, writing the Lagrangian as

$$(121) \mathcal{L} = -A_\alpha(x)\hat{M}_{\alpha\beta}A_\beta(x) - \bar{\psi}_\alpha(x)\hat{D}_{\alpha\beta}\psi_\beta(x) + \mathcal{L}^{\text{int}},$$

where \mathcal{L}^{int} contains all trilinear and quadrilinear terms. Here, $A_\alpha(x)$ is short for all bosonic (scalar and vector) fields, and $\bar{\psi}$ and ψ for both the Dirac fermions and the Faddeev-Popov fermions. The coefficients $\hat{M}_{\alpha\beta}$, $\hat{D}_{\alpha\beta}$ and the trilinear coefficients may contain the gradient operator $\partial/\partial x^\mu$. After Fourier expansion, this will turn into a factor ik_μ .

- The propagators $\hat{P}_{\alpha\beta}$ and $\hat{P}_{\alpha\beta}^{\text{ferm}}$ will be the inverse of the coefficients $\hat{M} - i\varepsilon$ and $\hat{D} - i\varepsilon$, so, for instance

$$\begin{aligned} \text{if } \hat{M}_{\alpha\beta} &= (m_{(\alpha)} - \partial_\mu^2)\delta_{\alpha\beta} \quad \text{then} \quad \hat{P}_{\alpha\beta} = \frac{\delta_{\alpha\beta}}{m_{(\alpha)}^2 + k^2 - i\varepsilon}; \\ \text{if } \hat{D}_{\alpha\beta} &= (m_{(\alpha)} + \gamma^\mu\partial_\mu)\delta_{\alpha\beta} \quad \text{then} \quad \hat{P}_{\alpha\beta}^{\text{ferm}} = \frac{(m_{(\alpha)} - i\gamma^\mu k_\mu)\delta_{\alpha\beta}}{m_{(\alpha)}^2 + k^2 - i\varepsilon}. \end{aligned} \quad (122)$$

- The vertices are generated by the trilinear and quadrilinear terms of \mathcal{L}^{int} , just as in subsection 2.5. If we have source terms such as $J^a(x)\phi_a(x)$, $\bar{\eta}^i(x)\psi_i(x)$ or $\bar{\psi}^i(x)\eta_i(x)$, then these correspond to propagators ending into points, where the momentum k has to match a given Fourier component of the source. All this can be read off neatly from formal expansions of the functional integral such as (47).
- There is an overall minus sign for every fermionic closed loop.
- Every diagram comes with canonical coefficients such as $1/k!$ and $(2\pi)^{-4N}$ where $k!$ is the dimension of the diagram's internal symmetry group, and N counts the number of loop integrations. These coefficients can be obtained by comparing functional integrals with ordinary integrals.
- There is a normalization coefficient for every external line, depending on the wave function chosen for the in- and out-going particles. We return to this in section 6.

⁶One may decide to leave small corrections to the bilinear parts of the Lagrangian to be treated together with the higher order terms as if they were 'two-point vertices'.

Note that any terms in the Lagrangian that can be written as a gradient of some (locally defined) field configuration can be replaced by zero. This is because (under sufficiently carefully chosen boundary conditions) such terms do not contribute to the total action $S = \int d^4x \mathcal{L}(x)$.

4.6 BRST symmetry

As the reader may have noted, we departed from our original intention, to keep space and time on a lattice and only turn to the continuum limit at the very end of a calculation. We have not even started doing calculations, and already the Feynman rules were formulated as if the fields lived on a space-time continuum. Indeed, we should have kept space and time discrete, so that the functional integral is nothing but an ordinary integral in a space with very many, but still a finite number of, dimensions. In practice, however, the continuum is a lot easier to handle, so, often we do not explicitly mention the finite size meshes of space and time.

Our first attempt to formulate the continuum limit will be in section 7. We will then see that the coefficients in the Lagrangian (120) have to be *renormalized*. The following question then comes up:

If we see a Lagrangian that looks like (120), how can we check that its coefficients are those of a genuine gauge theory?

The answer to this question is that the gauge-fixed Lagrangians (4.35) and (4.36) possess a *symmetry*. The first attempts to identify the symmetry in question gave negative results, because the ghost field is fermionic while the gauge fixing terms are bosonic. In the early days we thought that the required relation between the gauge fixing terms and the ghost terms had to be checked by inspection [’t Hooft and Veltman, 1972a]. But the complete answer was discovered by Becchi, Rouet and Stora [1975; 1976], and independently by Tyutin [Tyutin, 1975]. The symmetry, called BRST symmetry, is a *supersymmetry*. For the Lagrangian (120), which is slightly more general than (119), the transformation rules are

$$\begin{aligned}
 A'_\alpha(x) &= A_\alpha(x) + \bar{\varepsilon} \frac{\partial A_\alpha(x)}{\partial \Lambda^b(x')} \eta^b(x'); & (a) \\
 \eta^{a'}(x) &= \eta^a(x) + \frac{1}{2} \bar{\varepsilon} f_{abc} \eta^b(x) \eta^c(x); & (b) \\
 \bar{\eta}^{a'}(x) &= \bar{\eta}^a(x) + \bar{\varepsilon} C^a(x), & (c)
 \end{aligned}
 \tag{123}$$

where the anticommuting number $\bar{\varepsilon}$ is the infinitesimal generator of this (global) supersymmetry transformation.

The invariance of the Lagrangian (120) under this supersymmetry transformation is easy to check, except perhaps the cancellation of the variation of the last

term against the contribution of (123 b):

$$(124) \quad \bar{\eta}^a \frac{\partial C^a}{\partial \Lambda^b} \frac{1}{2} \bar{\varepsilon} f_{bcd} \eta^c \eta^d + \bar{\eta}^a \gamma_b \frac{\partial}{\partial \Lambda^c} \frac{\partial C^a}{\partial \Lambda^d} \eta^c \eta^d = \dots$$

Substituting some practical examples for the gauge constraint function C^a , one discovers that these terms always cancel out. The reason for (124) to vanish is the fact that gauge transformations form a group, implying the Jacobi identity:

$$(125) \quad f_{asb} f_{scd} + f_{asc} f_{sdb} + f_{asd} f_{sbc} = 0.$$

The converse is more difficult to prove: *If* a theory is invariant under a transformation of the form (123) (BRST invariance), then it is a gauge-fixed local gauge theory. What is really needed in practice, is to show that the ghost particles do not contribute to the S -matrix. This indeed follows from BRST invariance, via the so-called Slavnov-Taylor identities [Slavnov, 1972; Taylor, 1971], relations between amplitudes that follow from this symmetry.

5 THE BROUT-ENGLERT-HIGGS MECHANISM

The way it is described above, Yang-Mills gauge theory does not appear to be suitable to describe *massive* particles with spin one. However, in our approach we concentrated only on the high-energy, high-momentum limit of theories for vector particles, by assuming the Lagrangian to take the form (85) there. Mass terms dominate in the infra-red, or low energy domain. Here, one may note that we have not yet exploited all possibilities.

We need to impose exact local gauge-invariance, as explained in subsection 4.3. So our theory must be constructed along the lines expounded in subsection 4.2. All scalar and spinor fields must come as representations of the gauge group. So, what did we overlook?

In our description of the most general, locally gauge-invariant Lagrangian, it was tacitly assumed that the minimum of the scalar potential function $V(\phi)$ occurs at $\phi = 0$, so that, as one may have in global symmetries, the symmetry is evident in the particle spectrum: physical particles come as representations of the full local symmetry group. But, as we have seen in the case of a global symmetry, in subsection 2.2, the minimum of the potential may occur at other values of ϕ . If these values are not invariant under the gauge group, then they form a non-trivial representation of the group, invariant only under a subgroup of the gauge group. It is the invariant subgroup, if at all non-trivial, of which the physical particles will form representations, but the rest of the symmetry is hidden. Indeed, if we switch off the coupling to the vector fields, we obtain again the situation described in subsection 2.2. As was emphasized there, the particle spectrum then contains

massless particles, the Goldstone bosons. These Goldstone bosons represent the field excitations associated to a global symmetry transformation, which does not affect the energy: hence the absence of mass.

But, global gauge Goldstone bosons *do* carry a kinetic term. Therefore, they do carry away energy when moving with the speed of light. This is because a global symmetry only dictates the Goldstone field to carry no energy if the field is space-time independent.

In contrast, *local* gauge symmetries demand that Goldstone fields also carry no energy when they do depend on space and time. In the case of a local symmetry, therefore, Goldstone modes are entirely in the ghost sector of the theory; Goldstone particles then are unphysical. Let us see how this happens in an example.

5.1 The $SO(3)$ case

As a prototype, we take the group $SO(3)$ as our local gauge group, and for simplicity we ignore the contributions of loop diagrams, which represent the higher order quantum corrections to the field equations. Let the scalar field ϕ_a be in the 3-representation. The invariant part of the Lagrangian is then:

$$(126) \quad \mathcal{L}^{\text{inv}} = -\frac{1}{4}(F_{\mu\nu}^a)^2 - \frac{1}{2}(D_\mu\phi_a)^2 - V(\phi); \quad V(\phi) = \frac{1}{8}\lambda((\phi_a)^2 - F^2)^2.$$

Here, D_μ stands for the covariant derivative: $D_\mu\phi_a = \partial_\mu\phi_a + g\epsilon_{abc}A_\mu^b\phi_c$. As in section 2.2, Eq. (18), we define shifted fields $\tilde{\phi}_a$ by

$$(127) \quad \phi_a \equiv \tilde{\phi}_a + \begin{pmatrix} 0 \\ 0 \\ F \end{pmatrix}; \quad V(\tilde{\phi}) = \frac{1}{2}\lambda F^2\tilde{\phi}_3^2 + \frac{1}{2}\lambda F\tilde{\phi}^2\tilde{\phi}_3 + \frac{1}{8}\lambda(\tilde{\phi}^2)^2.$$

The shift must also be carried out in the kinetic term for ϕ :

$$(128) \quad \begin{aligned} D_\mu\phi_a &= D_\mu\tilde{\phi}_a + gF \begin{pmatrix} A_\mu^2 \\ -A_\mu^1 \end{pmatrix}; \quad -\frac{1}{2}(D_\mu\phi_a)^2 = \\ &-\frac{1}{2}(D_\mu\tilde{\phi}_a)^2 - gF(A_\mu^2 D_\mu\tilde{\phi}_1 - A_\mu^1 D_\mu\tilde{\phi}_2) - \frac{1}{2}g^2 F^2(A_\mu^1{}^2 + A_\mu^2{}^2). \end{aligned}$$

Defining the complex fields

$$(129) \quad \begin{aligned} \tilde{\Phi} &= \frac{1}{\sqrt{2}}(\tilde{\phi}_1 + i\tilde{\phi}_2); \quad \mathcal{A}_\mu = \frac{1}{\sqrt{2}}(A_\mu^1 + iA_\mu^2); \\ D_\mu\tilde{\Phi} &= (\partial_\mu + iA_\mu^3)\tilde{\Phi} - i\mathcal{A}_\mu\tilde{\phi}_3, \end{aligned}$$

we see that the Lagrangian (126) becomes

$$\begin{aligned} \mathcal{L}^{\text{inv}} &= -\frac{1}{4}(F_{\mu\nu}^a)^2 - \frac{1}{2}(D_\mu \tilde{\phi}_3)^2 - D_\mu \tilde{\Phi}^* D_\mu \tilde{\Phi} \\ &\quad - \frac{1}{2} M_H^2 \tilde{\phi}_3^2 - M_V^2 \mathcal{A}_\mu^* \mathcal{A}_\mu + M_V \Im(\mathcal{A}_\mu^* D_\mu \tilde{\Phi}) - V^{\text{int}}(\tilde{\phi}), \end{aligned}$$

(130) where $M_H = \sqrt{\lambda}F$; $M_V = gF$,

and V^{int} is the remainder of the potential term. \Im stands for imaginary part.

Thus, the ‘neutral’ component of the scalar field, the Higgs particle, gets a mass M_H (see Eq. 127) and the ‘charged’ components of the vector field receive a mass term with mass M_V . The mechanism that removes (some of) the Goldstone bosons and generates mass for the vector particles, is called the Brout-Englert-Higgs (BEH) mechanism [Englert and Brout, 1964; Higgs, 1964b; Higgs, 1964a; Higgs, 1966]. In every respect, the neutral, massless component of the vector field behaves like an electromagnetic vector potential, and the complex vector particle is electrically charged.

5.2 Fixing the gauge

If one would try to use the rules of Subsection 4.5 to derive the Feynman rules directly from \mathcal{L}^{inv} , one would find that the matrix \hat{M} describing the bilinear part of the Lagrangian has no inverse. This is because the gauge must first be fixed. Choosing $\partial_\mu A_\mu^a(x) = 0$ has the advantage that the somewhat awkward term $\Im(\mathcal{A}_\mu^* \partial_\mu \tilde{\Phi})$ can be put equal to zero by partial integration. The vector propagator (in momentum space) is then easily computed to be

$$(131) P_{\mu\nu}^{ab}(k) = \frac{\delta_{\mu\nu} - k_\mu k_\nu / (k^2 - i\varepsilon)}{k^2 + m_{(a)}^2 - i\varepsilon} \delta_{ab},$$

where $m_{(a)} = M_V$ for the charged vector field and 0 for the neutral one.

This indeed appears to describe a vector particle with mass $m_{(a)}$ and an additional transversality constraint. One can do something smarter, though. If, in the gauge-fixed lagrangian (120), we choose

$$(132) C^3 = \partial_\mu A_\mu^3; \quad C^1 = \partial_\mu A_\mu^1 - M_V \tilde{\phi}_2; \quad C^2 = \partial_\mu A_\mu^2 + M_V \tilde{\phi}_1,$$

then we find that the scalar-vector mixing terms cancel out, but now also the $(\partial_\mu A_\mu)^2$ term cancels out, so that the vector propagator loses its $k_\mu k_\nu$ term. The vector propagator is then

$$(133) P_{\mu\nu}^{ab}(k) = \frac{\delta_{\mu\nu} \delta_{ab}}{k^2 + m_{(a)}^2 - i\varepsilon},$$

and the charged scalar ghost gets a mass M_V . The physical field $\tilde{\phi}_3$ is unaffected.

It is instructive to compute the Faddeev-Popov ghost Lagrangian in this gauge. One easily finds it to be

$$(134) \mathcal{L}^{\text{ghost}} = \bar{\eta}^a \partial^2 \eta^a - M_V^2 (\bar{\eta}^1 \eta^1 + \bar{\eta}^2 \eta^2) + \text{interaction terms}.$$

As can be confirmed by more explicit calculations, the theory has physical, charged vector particles with masses M_V , a neutral (massless) photon and a neutral scalar particle with mass M_H . The latter is called the Higgs particle of this theory. All other fields in the Lagrangian describe ghost fields. Apparently, in the gauge described above, all ‘unphysical’ charged particles, the ghosts, the timelike components of the vector fields, as well as the Goldstone bosons, have the same mass M_V . The unphysical neutral particles all have mass zero.

One concludes that the symmetry pattern of this example is as follows: the local gauge group, $SO(3)$, is broken by the Brout-Englert-Higgs mechanism into its subgroup $SO(2)$ (the rotations about a fixed axis, formed by the vacuum value of ϕ_a), or equivalently, $U(1)$. Therefore two of the three vector bosons obtain a mass, while one massless $U(1)$ photon remains. At the same time, two of the three scalars turn into ghosts, the third into a Higgs particle.

The Brout-Englert-Higgs mechanism does not alter the total number of independent physical states in the particle spectrum. In our example, two of the three scalar particles disappeared, but the two massive spin-1 particles now each have three spin helicities, whereas the massless photons only had two.

5.3 Coupling to other fields

The shift (127) in the definition of the fields, gives all interactions an asymmetric appearance. This is why, in the literature, one talks of “spontaneous breaking of the local symmetry”. Actually, this is something of a misnomer. In the case of a *global* symmetry, spontaneous breakdown means that the vacuum state is degenerate. After a global symmetry transformation, the vacuum state is transformed into a physically inequivalent vacuum state, which is not realized in the system. The existence of a massless Goldstone boson testifies to that. In the case of a *local* symmetry, nothing of the sort happens. There is only one vacuum state, and it is invariant under the local symmetry, always. This is why the Goldstone boson became unphysical. In fact, *all* physical states are formally invariant under local gauge transformations. Apparent exceptions to this rule are, of course, the charged particles in QED, but this is because we usually wish to ignore their interactions with the vector potential at infinity. In reality a full discussion of charged particles is obscured by their long-range interactions.

In view of all of this, it is better not to say that a local symmetry is spontaneously broken, but, rather, to talk of the *Brout-Englert-Higgs mechanism* [Englert

and Brout, 1964; Higgs, 1964b], which is the phenomenon that the spectrum of physical particles do not form a representation of the local symmetry group. The local symmetry can only be recognized by shifting the scalar fields back to their symmetric notation, the original fields ϕ . Local symmetry must not be regarded as a property of the physical states, but rather as a property of our way of describing the physical states.

If, however, we perform a perturbation expansion for small values of the gauge field coupling, we find that at vanishing gauge coupling a local symmetry *is* spontaneously broken. Therefore, it is still quite useful to characterize our perturbative description by listing the gauge groups and the subgroups into which they are broken.

Now, let us assume that there are other fields present, such as the Dirac fermions, ψ_i . In the symmetric notation, they must form a representation of the local gauge group. So, we have

$$(135) \mathcal{L}^{\text{Dirac}} = -\bar{\psi}^i (\gamma_\mu D_\mu + m_{(i)}) \psi_i - \bar{\psi}^i g_Y t_{ij}^a \phi_a \psi_j,$$

where D_μ is the appropriate covariant derivative, containing those matrices T^a that are appropriate for the given representation (see 95 and 96), and g_Y stands for one or more Yukawa coupling parameters. The mass terms $m_{(i)}$ and coupling coefficients t_{ij}^a are invariant tensors of the gauge group (masses are only allowed if the fermions are not chiral, see the discussion following Eq. (106)).

Here, again, we started with the more transparent symmetric fields ϕ_a , but the physical fields $\tilde{\phi}$ are obtained by the shift $\phi_a = F_a + \tilde{\phi}_a$. Thus, the lowest order bilinear part of the Dirac Lagrangian becomes

$$(136) \mathcal{L}^{\text{Dirac}} \rightarrow -\bar{\psi}^i \left((\gamma_\mu \partial_\mu + m_{(i)}) \delta_{ij} + g_Y t_{ij}^a F_a \right) \psi_j,$$

In particular, if the symmetry acts distinctly on the chiral parts of the fermion fields, the mass term $m_{(i)}$ is forbidden, but the less symmetric second term may generate masses and in any case mass differences for the fermions. Thus, not only do the vector and scalar particles no longer form representations of the original local gauge group, but neither do the fermions.

5.4 The Standard Model

What is presently called the ‘Standard Model’ is just an example of a Higgs theory. The gauge group is $SU(3) \times SU(2) \times U(1)$. This means that the set of vector fields falls apart into three groups: 8 associated to $SU(3)$, then 3 for $SU(2)$, and finally one for $U(1)$. The scalar fields ϕ_i form one two-dimensional, complex representation of two of the three groups: it is a doublet under $SU(2)$ and rotates as a particle with charge $\frac{1}{2}$ under $U(1)$.

Representing the Higgs scalar in terms of four real field components, the Brout-Englert-Higgs mechanism is found to remove three of them, leaving only one neutral, physical Higgs particle. $SU(2) \times U(1)$ is broken into a diagonal subgroup $U(1)$. Three of the four gauge fields gain a mass. The one surviving photon field is obtained after re-diagonalizing the vector fields; it is a linear composition of the original $U(1)$ field and one of the three components of the $SU(2)$ gauge fields.

The $SU(3)$ group is not affected by the Brout-Englert-Higgs mechanism, so one would expect all ‘physical’ particles to come in representations of $SU(3)$. What happens instead is further explained in section 11: only gauge-invariant combinations of fields are observable as particles in our detectors.

The fermions in the Standard Model form three ‘families’. In each family, we see the same pattern. The left handed fields, ψ_L , all form doublets under $SU(2)$, and a combination of a triplet (‘quarks’) and singlets (‘leptons’) under $SU(3)$. The right handed components, ψ_R , form the same representations under $SU(3)$, but form a pair of two singlets under $SU(2)$; so they do not couple to the $SU(2)$ vector fields. The $U(1)$ charges of the left-handed $SU(2)$ doublets are $-\frac{1}{2}$ for the leptons and $\frac{1}{6}$ for the quarks; the $U(1)$ charges of the right-handed singlets are -1 and 0 (for the leptons) or $-\frac{1}{3}$ and $\frac{2}{3}$ (for the quarks).

The Standard Model owes its structure to the various possible Yukawa interaction terms with the Higgs scalars. They are all of the form $\bar{\psi} \phi \psi$, and invariant under the entire gauge group, but since there are three families of fermions, each having left and right handed chiral components, there are still a fairly large number of such terms, each of which describes an interaction strength whose value is not dictated by the principles of our theory [Hoddeson *et al.*, 1997].

6 UNITARITY

As we saw in subsection 4.5, the Feynman rules unambiguously follow from the expression one has for the Lagrangian of the theory. More precisely, what was derived there was the set of rules for the vacuum-to-vacuum amplitude in the presence of possible source insertions $J_i(x)$, including anticommuting sources $\eta_i, \bar{\eta}_i$. The overall multiplicative constant C in our Gaussian integrals such as (47) is completely fixed by the demand that, in the absence of sources, the vacuum-to-vacuum amplitude should be 1. By construction then, the resulting scattering matrix should turn out to be unitary.

In practice, however, things are not quite that simple. In actual calculations, one often encounters divergent, hence meaningless expressions. This happens when one makes the transition to the continuum limit too soon — remember that we insisted that space and time are first kept discrete. Unitarity of the S -matrix turns out to be a sensitive criterion to check whether we are performing the continuum limit correctly. It was one of our primary demands when we initiated the

program of constructing workable models for relativistic, quantized particles. Another demand, the validity of dispersion relations, can be handled the same way as unitarity; these two concepts will be shown to be closely related. The formalism described below is based on work by Cutkosky and others, but was greatly simplified by Veltman [’t Hooft and Veltman, 1994].

Parts of this section are fairly technical and could be skipped at first reading.

6.1 The largest time equation

Let us start with the elementary Feynman propagator, $(k^2 + m^2 - i\varepsilon)^{-1}$, and its Fourier transform back to configuration space (omitting for simplicity a factor $(2\pi)^4$):

$$(137) \Delta^F(x) = -i \int d^4k \frac{e^{ikx}}{k^2 + m^2 - i\varepsilon}, \quad x = x^{(1)} - x^{(2)}.$$

In addition, we define the *on-shell propagators*

$$(138) \Delta^\pm(x) = 2\pi \int d^4k e^{ikx} \delta(k^2 + m^2) \theta(\pm k^0); \quad kx = \vec{k} \cdot \vec{x} - k^0 x^0,$$

and θ is the Heaviside step function, $\theta(x) = 1$ for $x \geq 0$ and $= 0$ otherwise. The integrals are over Minkowski variables \vec{k}, k^0 . The operators (138) propagate particles on mass shell with the given sign of the energy from $x^{(2)}$ to $x^{(1)}$, or back with the opposite sign. We have

$$(139) \Delta^+(x) = (\Delta^-(x))^*; \quad \Delta^+(x) = \Delta^-(-x).$$

Our starting point is the decomposition of the propagator into forward and backward parts:

$$(140) \Delta^F(x) = \theta(x^0) \Delta^+(x) + \theta(-x^0) \Delta^-(x).$$

Obviously:

$$(141) \Delta^{F*}(x) = \theta(x^0) \Delta^-(x) + \theta(-x^0) \Delta^+(x).$$

One easily proves this by deforming the contour integration in the complex k^0 plane.

Consider now a Feynman diagram with n vertices, where lines are attached with a given topological structure, which will be kept fixed. The external lines are assumed to be ‘amputated’: there are no propagators attached to them. The

Feynman rules are applied as described in Subsections 2.5 and 4.5. The diagram is then part of our calculation of an S -matrix element. We consider the diagram in momentum representation *and* in the coordinate representation. The expression we get in coordinate representation is called $F(x^{(1)}, x^{(2)}, \dots, x^{(n)})$.

Next, we introduce an expression associated to the same diagram, but where some of the vertices are underlined:

$$F(x^{(1)}, x^{(2)}, \dots, \underline{x}^{(i)}, \dots, \underline{x}^{(j)}, \dots, x^{(n)}),$$

where $x^{(i)}$ refer to the coordinates that must be integrated over when one elaborates the Feynman rules. The rules for computing this new amplitude are as follows:

- i*) A propagator $\Delta^F(x^{(i)} - x^{(j)})$ is used if neither $x^{(i)}$ nor $x^{(j)}$ are underlined.
- ii*) A propagator $\Delta^+(\underline{x}^{(i)} - x^{(j)})$ is used if $x^{(i)}$ but not $x^{(j)}$ is underlined.
- iii*) A propagator $\Delta^-(x^{(i)} - \underline{x}^{(j)})$ is used if $x^{(j)}$ but not $x^{(i)}$ is underlined.
- iv*) A propagator $\Delta^{F*}(\underline{x}^{(i)} - \underline{x}^{(j)})$ is used if both $x^{(i)}$ and $x^{(j)}$ are underlined.
- v*) A minus sign is added for every underlined vertex.

In all other respects, the rules for the calculation of the amplitude are unchanged.




Figure 3. Diagram with underlined vertices, which are indicated by little circles

One now derives the *largest time equation*:

Let $x^{(k)}$ be the coordinate with the largest time:

$$x^{(k)0} \geq x^{(i)0}, \forall i.$$

Then,

$$(142) F(x^{(1)}, x^{(2)}, \dots, x^{(k)}, \dots, \underline{x}^{(n)}) = -F(x^{(1)}, x^{(2)}, \dots, \underline{x}^{(k)}, \dots, \underline{x}^{(n)}),$$

where in both terms the points other than $x^{(k)}$ are underlined or not in identical ways.

One easily proves this using Eqs. (140) and (141). One consequence of this theorem is

$$(143) \sum_{\text{all } 2^n \text{ possible underlinings}} F(\{x^{(i)}\}) = 0.$$

We now show that these are the diagrams contributing to the unitarity equation, or ‘optical theorem’:

$$(144) \sum_n S|n\rangle\langle n|S^\dagger = \mathbf{1}.$$

The diagrams for the matrix S are as described earlier. The diagrams for S^\dagger contain the complex conjugates of the propagators. Since also the vertices in the functional integral are all multiplied by i , they must all change sign in S^\dagger . Also the momenta k in e^{ikx} switch sign. In short, the diagrams needed for the computation of S^\dagger indeed are the underlined Green functions. Note that, in momentum space, the largest time equation (142) cannot be applied to individual vertices, since, while being integrated over, the vertex with largest time switches position. However, the summed equation (143) is valid. The identity $\mathbf{1}$ on the r.h.s. of Eq. (144) comes from the one structure that survives: the diagram with no vertices at all.

We observe that unitarity may follow if we add all possible ways in which a diagram with given topology can be cut in two, as depicted in Fig. 3. The shaded line separates S from S^\dagger .

The lines joining S to S^\dagger represent the intermediate states $|n\rangle$ in Eq. (144). They are on mass shell and have positive energies, which is why we need the factors $\delta(k^2 + m^2)\theta(k^0)$ there. If a propagator is equipped with some extra coefficients R_{ij} :

$$(145) P_{ij}(k) = \frac{-i R_{ij}(k)}{k^2 + m^2 - i\varepsilon},$$

then we can still use the same decomposition (140), provided R_{ij} is *local*: it must be a finite polynomial in k . Writing

$$(146) R_{ij} = \sum_k f_i(k) f_j^*(k),$$

we can absorb the factors $f_i(k)$ into the definition of S , provided that all eigenvalues of R_{ij} are non-negative. Indeed, kinetic terms in the Lagrangian must all have the same sign.

Note that we are *not* allowed to replace the terms in the Lagrangian by their complex conjugates. This implies that, for the unitarity proof, it is mandatory that the Lagrange density is a real function of the fields.

An important feature of these equations is the theta functions for k^0 . They guarantee that the intermediate states contribute only if their total energy does not exceed the energy available in the given channel.

6.2 Dressed propagators

In the previous subsection, not all diagrams that contribute to $S S^\dagger$ have yet been handled correctly. There is a complication when self-energy diagrams occur. If one of the lines at both sides of a self-energy blob is replaced by Δ^\pm , then the other propagator Δ^F places a pole on top of that Dirac delta. In this case, we have to use a more sophisticated prescription. To see what happens, we must first sum the geometric series of propagator insertions, see Fig. 4(a). We obtain what is called the dressed propagator. In momentum space, let us write the contribution of a single blob in Fig. 4(a) as $-i\delta M(k)$. It represents the summed contribution of all *irreducible* diagrams, which are the diagrams with two external lines that cannot fall apart if one cuts one internal line. We need its real and imaginary parts: $\delta M(k) \equiv \delta m^2(k) - i\Gamma(k)$. Write the full propagator as

$$\begin{aligned} P^{\text{dr}}(k) &= P^0(k) - P^0(k)i\delta M(k)P^0(k) + \dots \\ &= P^0(k) \sum_{n=0}^{\infty} \left(-i\delta M(k)P^0(k) \right)^n \\ &= \frac{P^0(k)}{1 + i\delta M(k)P^0(k)}; \end{aligned} \tag{147}$$

$$\begin{aligned} \text{if } P^0(k) &= -i(M(k) - i\varepsilon)^{-1} \\ \text{then } P^{\text{dr}}(k) &= -i(M(k) + \delta M(k) - i\varepsilon)^{-1}, \end{aligned} \tag{148}$$

where $P^0(k)$ is the unperturbed ('bare') propagator.

If we define the real part of the dressed propagator (in momentum space) to be

$$(149) \quad \Re(P^{\text{dr}}(k)) = \frac{\Gamma(k)}{(k^2 + M + \delta m^2)^2 + \Gamma^2} = \pi \varrho(-k^2),$$

then, by contour integration,

$$(150) \quad P^{\text{dr}}(k) = \int_0^\infty dm^2 \frac{\varrho(m^2)}{k^2 + m^2 - i\varepsilon};$$

we call this the Källén-Lehmann representation of the propagator. Later, it will be assured that $\varrho(m^2) = 0$ if $m^2 < 0$.




Figure 4. (a) The dressed propagator as a geometric series; (b) Cutting the dressed propagator

The best strategy now is to apply a largest time equation to the entire dressed propagator. Write, instead of Eqs. (140) and (141),

$$\begin{aligned}
 P^{\text{dr}}(x) &= \theta(x^0)\Delta_{\text{dr}}^+(x) + \theta(-x^0)\Delta_{\text{dr}}^-(x); \\
 (151) \quad P^{\text{dr}}(x)^* &= \theta(x^0)\Delta_{\text{dr}}^-(x) + \theta(-x^0)\Delta_{\text{dr}}^+(x).
 \end{aligned}$$

Then,

$$(152) \quad \Delta_{\text{dr}}^\pm(k) = 2\pi \int d^4k e^{ikx} \varrho(-k^2)\theta(\pm k^0).$$

The imaginary part $\Gamma(k)$ of the irreducible diagrams can itself again be found by applying the cutting rules. Writing $S = \mathbf{I} + iT$, we find that unitarity for all non-trivial diagrams corresponds to $i(T - T^\dagger) + TT^\dagger = 0$, and the diagrams for TT^\dagger are depicted in Fig. 4b. They are exactly the diagrams needed for unitarity of the entire scattering matrix involving a single virtual particle in the channel of two external ones.

One observes that the function $\varrho(-k^2)$ must be non-negative, and only nonvanishing for timelike k . The latter is guaranteed by the theta functions in k^0 . Only the delta peaks in ϱ are associated to *stable* particles that occur in the initial and final states of the scattering matrix. Resonances with finite widths contribute to the unitarity of the scattering matrix via their stable decay products.

6.3 Wave functions for in- and out-going particles

Many technical details would require too much space for a full discussion, so we have to keep this sketchy. In case we are dealing with vector or spinor particles, the residues R_{ij} of the propagators represent the summed absolute squares of the particle wave functions. We have seen in Eq. (145) how this comes about. If, for

example, a vector particle is described by a propagator

$$(153) P_{\mu\nu} = -i \frac{\delta_{\mu\nu} + k_\mu k_\nu / M^2}{k^2 + M^2 - i\varepsilon},$$

then we see that, first of all, the numerator is a polynomial in k , as was required, and, if we go on mass shell, $k^2 = -M^2$, then we see that the field component proportional to k_μ is projected out. In particular, if we put $k = (0, 0, 0, iM)$, then $R_{ij} = \delta_{ij}$ and its timelike components disappear, so indeed there are three independent states for the particle described.

For the fermions, the bare propagator is

$$(154) P^{\text{Dirac}} = -i \frac{m - i\gamma k}{k^2 + m^2 - i\varepsilon}.$$

Before relating this to the renormalization of the wave functions, we must note that all γ^μ are hermitean, while k_i are real and k_4 is imaginary. We observe that the Feynman rules for S^\dagger are like those of S , but with γ^4 replaced by $-\gamma^4$. Next, the arrows in the propagators must be reversed. This leads to an extra minus sign for every vector k_μ , while γ^μ are replaced by $\gamma^{\mu\dagger}$. All together, one requires that $\gamma^i \rightarrow -\gamma^i$ while γ^4 remains unchanged. This amounts to the replacement $\gamma^\mu \rightarrow \gamma^4 \gamma^\mu \gamma^4$. One concludes that the rules for S^\dagger are like those for S if all fermion lines enter or leave the diagram with an extra factor γ^4 . This means that the wave functions for external fermions in a diagram are to be normalized as

$$(155) (m - i\gamma_i k^i + \gamma^4 k^0) \gamma^4 = \sum_{i=1}^2 |\psi_i(k)\rangle \langle \psi_i(k)|, \quad (k^0 > 0);$$

while for the anti-fermions, we must demand

$$(156) \gamma^4 (m - i\gamma_i k^i - \gamma^4 k^0) = - \sum_{i=3}^4 |\psi_i(k)\rangle \langle \psi_i(k)|, \quad (k^0 > 0).$$

The minus sign is necessary because the operator in (156) has two negative eigenvalues. One concludes that unitarity requires spin- $\frac{1}{2}$ particles to carry one extra minus sign for each closed loop of these particles. This leads to the necessity of Fermi-Dirac statistics. Again, it is important that none of the higher order corrections ever affect the signs of the eigenvalues for these projection operators, since these can never be accommodated for by a renormalization of the particle wave function.

The conclusion of this section, that the scattering matrix is unitary in the space of physical particle states, should not come as a surprise because our theory has

been constructed to be that way. Yet it is important that we see here in what way the Feynman diagrams intertwine to produce unitarity explicitly.

We also see that unitarity is much more difficult to control when we have ghosts due to the gauge fixing procedure. Our vector particles then have propagators where Eq. (153) is replaced by expressions such as

$$(157) P_{\mu\nu}^{\text{ren}} = \frac{-i g_{\mu\nu}}{k^2 + M^2 - i\varepsilon}.$$

We write here $g_{\mu\nu}$ rather than $\delta_{\mu\nu}$ in order to emphasize that our arguments are applied in Minkowski space, where clearly the time components ‘carry the wrong sign’. The field components associated to that would correspond to particles that contribute negatively to the scattering probabilities. To correct this, one would have to replace $|n\rangle\langle n|$ by $-|n\rangle\langle n|$, which cannot be achieved by renormalizing the states $|n\rangle$. Here, we use the BRST relations to show that all unphysical states can be transformed away. In practice, we use the fact that the scattering matrix does not depend on the choice of the gauge fixing function $C^a(x)$, so we choose it such that all ghost particles have a mass exceeding some critical value Λ . *In the intermediate states, their projector operators $\Delta^\pm(k)$ then only contribute if the total energy in the given channel exceeds Λ .* This then means that there are no ghosts in the intermediate states, so the scattering matrix is unitary in the space of physical particles only — an absolutely essential step in the argument that these theories are internally consistent. The required gauge fixing functions $C^a(x)$ are not difficult to construct, but their existence is only needed to complete this formal argument. They are rather cumbersome to use in practical calculations.

6.4 Dispersion relations

The largest time equation can also be employed to derive very important *dispersion relations* for the diagrams. These imply that any diagram D can be regarded as a combination of two sets of diagrams D_i and D_i^\dagger :

$$(158) D = \sum_i \int_0^\infty \frac{dk^0}{-k^0 - i\varepsilon} D_i(k^0) D_i^\dagger(k^0).$$

Here, $D_i(k^0)$ and $D_i^\dagger(k^0)$ stand for amplitudes depending on various external momenta k , where one of the timelike components, k^0 , is integrated over. This, one derives by singling out two points, $x^{(1)}$ and $x^{(2)}$ in a diagram, and time-ordering them. The details of the derivation go beyond the scope of this paper (although they are not more difficult than the previous derivations in this section). Eq. (158) can be used to express diagrams with closed loops in terms of diagrams with fewer closed loops, and discuss the subtraction procedures needed for renormalization.

7 RENORMALIZATION

For a proper discussion of the renormalization concept, we must emphasize what our starting point was: first, replace the continuum of space by a dense lattice of points, and only at the very end of all calculations do we make an attempt to take the continuum limit. The path integral procedure, illuminated in subsection 2.4, implies that time, also, can be replaced by a lattice. In Fourier space, the space-time lattice leads to finite domains for the values of energies and momenta (the Brillouin zones), so that all ultra-violet divergences disappear. If we also wish to ensure the absence of infra-red divergences, we must replace the infinite volume of space and time by a finite box. This is often required if complications arise due to divergent contributions of soft virtual particles, typically photons. Nasty infra-red divergences occur in theories with confinement, to be discussed in section 11.

The instruments that we shall use for the ultra-violet divergences of a theory are as follows. We assume that all freely adjustable physical constants of the theory, referred to as the 'bare' parameters, such as the 'bare' mass and charge of a particle, should be carefully tuned to agree with observation, but the tuning process may depend critically on the mesh size a of the space-time lattice. Thus, while we vary a , we allow all bare parameters, λ say, in the theories to depend on a , often tending either to infinity or to zero as $a \rightarrow 0$. If this procedure is combined with a perturbation expansion, say in terms of a small coupling g , we expect to find that observable features depend minimally on a provided that the bare couplings $g(a)$ remain small in the limit $a \downarrow 0$.

This will be an important condition for our theories to make sense at all. How do we know whether $g(a)$ tends to zero or not? The simplest thing to look at, is the *dimensionality* of g . All parameters of a field theory have a dimension of a length to some power. These dimensions usually depend on the dimensionality n of space-time. The rules to compute them are easy to obtain:

- An action $S = \int d^n x \mathcal{L}(x)$ is dimensionless;
- The dimension of a Lagrange density \mathcal{L} is therefore $(\text{length})^{-n} = m^n$, where m is a mass.
- The dimension of the fields can be read off from the kinetic terms in the Lagrangian, because they contain no further parameters. A scalar field ϕ has dimension $m^{(n-2)/2}$, a fermionic field ψ has dimension $m^{(n-1)/2}$.
- A gauge coupling constant g has dimension $m^{(4-n)/2}$ and the coupling parameter λ in an interaction term of the form $\lambda \phi^k$ has dimension $m^{n+k-nk/2}$,

and so on.

A theory is called power-counting renormalizable, if all expansion parameters have mass-dimension positive or zero.

This is why, in 4 space-time dimensions, we cannot accept higher than quartic interactions among scalars. In practice, in 4 space-time dimensions, most expansion parameters have dimension zero. In Section 9, we will see that dimensionless coupling parameters nevertheless depend on the size of a , but only logarithmically:

$$(159) \lambda(a) \approx \lambda_0 + C\lambda_0^2 \log(a) + \text{higher orders.}$$

Regardless of whether this tends to zero or to infinity in the continuum limit, one finds that, in the continuum theory, the perturbative corrections to the bare parameters λ diverge. This is nothing to be alarmed about. However, if λ itself is also a small parameter in terms of which we wish to perform a perturbation expansion, then clearly trouble is to be expected if its bare value tends to infinity. Indeed, we shall argue that, in general, such theories are inconsistent.

There are two very important remarks to be made:

- Theories can be constructed where all couplings really tend to zero in the continuum limit. These theories are called *asymptotically free* (Section 9), and they allow for accurate approximations in the ultra-violet. It is generally believed that such theories can be defined in a completely unambiguous fashion through their perturbation expansions in the ultra-violet; in any case, they allow for very accurate calculation of all their physical properties. QCD is the prime example.
- If a theory is not asymptotically free, but has only small coupling parameters, the perturbation expansion formally diverges, and the continuum limit formally does not exist. But the first N terms of perturbation expansion do make sense, where $N = \mathcal{O}(1/g)$. This means that uncontrollable margins of error are exponentially small, of order $e^{-C/g}$ or e^{-C/g^2} , which in practice is much smaller than other uncertainties in the theory, so they are of hardly any practical consequence. Thus, in such a case, our theory *does* have intrinsic inaccuracies, but these are *exponentially suppressed*. In practice, such theories are still highly valuable. The Standard Model is an example.

A useful approach is to substitute all *numbers* in a theory by formal *series expansions*, where the expansion parameter, a factor common to all coupling parameters of the theory, is formally kept infinitesimal. In that case, all perturbation coefficients are uniquely defined, though one has little direct knowledge concerning the convergence or divergence of the expansions.

In both the cases mentioned above, our theories are defined from their perturbation expansion; clearly, the perturbation expansion is not only a convenient device

for calculations, it is an essential ingredient in our theories. Let us therefore study how renormalization works, order-by-order in perturbation theory.

In a connected diagram, let the number of external lines be E , the number of propagators be P , and let V_n be the number of vertices with n prongs. By drawing two dots on each propagator and one on each external line, one finds that the number of dots is

$$(160) \quad 2P + E = \sum_n nV_n = 3V_3 + 4V_4.$$

For tree diagrams (simply connected diagrams), one finds by induction, with V the number of vertices, $V = \sum_n V_n$,

$$(161) \quad V = P + 1.$$

A diagram with L closed loops (an L -fold connected diagram) turns into a tree by cutting away L propagators. Therefore, one has

$$(162) \quad P = V - 1 + L.$$

Combining Eqs. (160) and (162), one has

$$(163) \quad E + 2L - 2 = \sum_n (n - 2)V_n = V_3 + 2V_4.$$

Consequently, if every 3-vertex comes with a factor g and every 4-vertex with a factor λ , and if a diagram with a given number E of external lines, behaves as $g^{2n}\lambda^k$, it must have $L = n + k + 1 - \frac{1}{2}E$ closed loops. Perturbation expansion is therefore often regarded as an expansion in terms of the number of closed loops.

7.1 Regularization schemes

In a tree diagram, in momentum space, no integrations are needed to be done — the momentum flowing through every propagator is fixed by the momenta of the in- and out-going particles. But if there are L loops, one has to perform $4L$ integrations in momentum space. It is these integrations that often tend to diverge at large momenta.

Of course, these divergences are stopped if momentum space is cut off, as is the case in a finite lattice. However, since our lattice is not Lorentz-invariant and may lack other symmetries such as gauge-invariance, it is useful to find *other ways to modify our theory so that UV divergences disappear*. This is called ‘regularization’. We give two examples.

Pauli-Villars regularization

Assume that a propagator of the form shown is replaced as follows:

$$(164) \quad \frac{A(k)}{k^2 + m^2 - i\varepsilon} \rightarrow \sum_i e_i \frac{A(k)}{k^2 + \Lambda_i^2 - i\varepsilon}; \quad \sum_i e_i = 0, \quad \sum_i e_i \Lambda_i^2 = 0.$$

If we take $e_1 = 1$, $\Lambda_1 = m$, while all other Λ_i tend to $+\infty$, we get back the original propagator. With finite Λ_i , however, we can make all momentum integrations converge at infinity. Our theory is then finite. This is (a somewhat simplified version of) Pauli-Villars regularization.

However, the new propagators cannot describe ordinary particles. The ones with $e_i < 0$ contribute to the unitarity relation with the wrong sign! On the other hand, the $i\varepsilon$ prescription is as usual, so that these particles do carry positive energy. In any channel where the total energy is less than Λ_i , the 'Pauli-Villars ghosts' do not contribute to the unitarity relation at all. So, in a theory where we put a limit to the total energy considered, Pauli-Villars regularization is physically acceptable. In practice, we will try to send all ghost masses Λ_i to infinity.

Dimensional regularization

Dimensional regularization [’t Hooft and Veltman, 1972b] consists of formally performing all loop integrations in $4 - \varepsilon$ dimensions, where ε may be any (possibly complex) number. As long as ε is irrational, all integrations can be replaced by finite expressions following an unambiguous prescription, to be explained below. If $\varepsilon = 0$, one can also subtract the integrals, but the prescription is then often not unambiguous, so that *anomalies* might arise. This is why dimensional regularization will be particularly important whenever the emergence of anomalies is a problem one wishes to understand and control.

It is important to realize that also when $\varepsilon \neq 0$, integrals may be divergent, but that, for irrational ε , unambiguous subtractions may be made. This needs to be explained, but first, one needs to define what it *means* to have non-integral dimensions. Such a definition is only well understood within the frame of the perturbation-, or loop-, expansion. Consider an irreducible diagram with L loops and N external lines, where we keep the external momenta $p_{(1)}, \dots, p_{(N)}$ fixed. It is obvious from the construction of the theory that the integrand is a purely rational function in $L(4 - \varepsilon)$ variables. Observing that the external momenta span some $N - 1$ dimensional space, we now employ the fact that the integration in the remaining dimensions is rotationally invariant. There, we write the formula for the ℓ -dimensional (Euclidean) sphere of radius r as

$$(165) \quad \int d^\ell k \delta(k^2 - r^2) = \frac{\pi^{\ell/2}}{\Gamma(\ell/2)} r^{\ell-2}.$$

Here, Γ stands for the Euler gamma function, $\Gamma(z) = (z-1)!$ for integral z .

It is at this point where we can decide that this expression *defines* the integral for any, possibly complex, value for ℓ . It converges towards the usual values whenever ℓ happens to be a positive integer. After having used Eq. (165), one ends up with an integral over s variables k_μ of a function $f(k)$, where s is an integer, but $f(k)$ contains ε -dependent powers of polynomials in k .

Convergence or divergence of an integral can be read off from simple power counting arguments, and, at first sight, one sees hardly any improvement when ε is close to zero. However, what is achieved is that infra-red divergences ($k_\mu \rightarrow 0$) are separated from the ultra-violet divergences ($k_\mu \rightarrow \infty$), and this allows us to define the “finite parts” of the integrals unambiguously:

- All integrals $\int d^s k f(k)$ are replaced by functionals $I(\{f(k)\})$ that obey the same combinatorial rules as ordinary integrals:

$$\begin{aligned} I(\alpha f_1 + \beta f_2) &= \alpha I(f_1) + \beta I(f_2), \\ (166) \quad I(\{f(k+q)\}) &= I(\{f(k)\}), \end{aligned}$$

- $I(f) = \int d^s k f(k)$ if this converges.
- $I(f) = 0$ if $f(k) = (k^2)^p$ when $2p + s$ is *not an integer*.

This latter condition is usually fulfilled, if we started with ε not integer.

These rules are sufficient to replace any integral one encounters in a Feynman diagram by some finite expression. Note, however, that complications arise if one wants to use these rules when $2p + s$ is an integer, particularly when it is zero. In that case, the expression diverges in the ultra-violet *and* in the infra-red, so, in this case, it cannot be used to remove all divergences — it can only replace one by another. Consequently, our finite expressions tend to infinity as $\varepsilon \rightarrow 0$.

It is important to verify that dimensional regularization fully respects unitarity and the dispersion relations discussed above. Therefore, the ‘dimensionally regularized’ diagrams correspond to solutions of the dispersion relations and the unitarity relations, providing some ‘natural’ subtraction.

Equivalence of regularization schemes

The subtractions provided by the various regularization schemes discussed above, in general, are not the same. At any given order, they do all obey the same dispersion relations of the form (158). If we ask, which amplitudes can be added to one scheme to reproduce the other, or, what is the amplitude of the difference between the two schemes (after having eliminated these differences at the order where the subdiagrams $D_i(k^0)$ had been computed), we find the following. This

difference must be a Lorentz-covariant expression; and it can only come from the dimensionally regularized contributions of the unphysical Pauli-Villars ghosts in Eq. (158). Because of their large masses, only very large values of k^0 in this equation contribute. The p^0 -dependence then must reduce to being a polynomial one (p being the momenta of the fixed external lines), and because of Lorentz-invariance, the expression must be polynomial in all components of p_μ . This is exactly what can be achieved by putting a *counter term inside the bare Lagrangian of the theory*. This way, one derives that the various regulators differ from one another by different effective couplings in the bare Lagrangian.

It is then a question of taste which regulator one prefers. Since dimensional regularization often completely respects local gauge-invariance⁷, and also because it turned out to be very convenient and efficient in practice, one often prefers that. It should always be kept in mind, however, that dimensional regularization is something of a mathematical trick, and the physical expressions only make sense in the limit $\varepsilon \rightarrow 0$.

7.2 Renormalization of gauge theories

Using the results from the previous Sections, we decide to treat quantum field theories in general, and gauge theories in particular, as follows: first, we *regularize* the theory, by using a ‘lattice cut-off’, or a Pauli-Villars cut-off, or by turning towards $n = 4 - \varepsilon$ dimensions. All these procedures are characterized by a small parameter, such as ε , such that the physical theory is formally obtained in the limit $\varepsilon \rightarrow 0$. These procedures are all equivalent, in the sense that by adding local interaction terms to the Lagrangian, one can map the results of one scheme onto those of another. Subsequently, we *renormalize* the theory. This means that all parameters in the Lagrangian are modified by finite corrections, which however may diverge in the limit $\varepsilon \rightarrow 0$. If these counter terms have been chosen well, the theory may stay finite and well defined in this limit. In particular, we should have a unitary, causal theory.

Unitarity is only guaranteed if the theory is gauge-invariant. Therefore, one prefers regulator schemes that preserve gauge-invariance throughout. This is what dimensional regularization often does. In that case, the renormalization procedure respects BRST-invariance, see Subsect. 4.6.

8 ANOMALIES

The Sections that follow will (again) be too brief to form a complete text for learning Quantum Field Theory. Our aim is here to give a summary of the features

⁷Only in one case, there is a complication, namely, when there are Adler-Bell-Jackiw anomalies; see Section 8.

that are all extremely important to understand the general structure of relativistic Quantum Field Theories.

If, for a given theory, no obviously gauge-invariant regularization procedure appears to exist, this might be for a reason: such a theory might not be renormalizable at all. In principle, this could be checked, as follows. One may always decide to use a regularization procedure that does not respect the symmetries one wants, provided that the symmetry can be restored in the limit where the physically observable effects of the regulator go away, such as $\varepsilon \rightarrow 0$, or $\Lambda_i \rightarrow \infty$, $i > 0$. If a gauge-invariant regulator does exist, but it hasn't yet been explicitly constructed, then we know that it differs from any other regulator by a bunch of finite counter terms. To find such counter terms is not hard, in practice; just add all terms needed to restore BRST invariance of the amplitudes.

But, in case that regulator is not known, how can we then be sure that such terms exist at all? BRST invariance requires the validity of the Slavnov-Taylor identities, but they appear to overdetermine the subtraction terms. This is the way we originally phrased the problem in [’t Hooft, 1971]. In fact, indeed there may be a clash. If this happens, it is called an *anomaly* [Jackiw, 1995, Ch.1].

Actually, the incidence of such anomalies is limited, fortunately. This is because for most theories completely gauge-invariant regulator techniques were found. Dimensional regularization often works. The one case where it does not is when there are *chiral fermions*. Classically, one may separate any fermionic field into a left-handed and a right handed part, as was mentioned in Subsection 4.2:

$$(167) \quad \begin{aligned} \psi(x) &= P_+ \psi_L(x) + P_- \psi_R(x); \quad P_{\pm} = \frac{1}{2}(1 \pm \gamma^5) \quad ; \\ \gamma_5 &= \frac{1}{24} \varepsilon_{\mu\nu\alpha\beta} \gamma^\mu \gamma^\nu \gamma^\alpha \gamma^\beta . \end{aligned}$$

Indeed, since $(\gamma^5)^2 = 1$, the operators P_{\pm} are genuine projection operators: $P_{\pm}^2 = P_{\pm}$.

The left- and right sectors of the fermions, see Eq. (106), may be separately gauge-invariant, transforming differently under gauge transformations. This, however, requires γ^5 to anti-commute with all other γ^μ , $\mu = 1, \dots, n$. But, as we see from their definition, Eq. (167), γ^5 only anti-commutes with four of the γ^μ , not all n . This is why the contributions from the $-\varepsilon$ remaining dimensions will not be gauge-invariant.

It was discovered by Bell and Jackiw [Bell and Jackiw, 1969], and independently by Adler [Adler, 1969; Adler and Bardeen, 1969; Bardeen, 1969], that no local counter term exists that obeys all symmetry conditions and has the desired dimensionality; Bell and Jackiw tried to use unconventional regulators, but those turned out not to be admissible. The basic culprit is the *triangle diagram*, Fig. 5(a), representing the matrix element of the axial vector current $\bar{\psi} \gamma^\mu \gamma^5 \psi$ in the field of two photons, each being coupled to the vector current $\bar{\psi} \gamma^\alpha \psi$.

For simplicity, we assume here the fermions to be massless. Let us call this




Figure 5. (a) The anomalous triangle diagram. μ, α and β are the polarizations, k, p and $q = k - p$ are the external momenta. (b) An anomalous diagram in non-Abelian theories

amplitude then $\Gamma_\mu^{\alpha,\beta}(p, q)$. It is linearly divergent. Upon regularization, there are two counter terms, or subtraction terms, whose coefficients should be determined, in a correct combination with the finite parts of the amplitude. Limiting ourselves to the correct quantum numbers and dimensions, we find the two quantities,

$$(168) \quad \begin{aligned} \delta_1 \Gamma_\mu^{\alpha,\beta}(p, q) &= \varepsilon_{\mu\alpha\beta\gamma} p_\gamma; \\ \delta_2 \Gamma_\mu^{\alpha,\beta}(p, q) &= \varepsilon_{\mu\alpha\beta\gamma} q_\gamma. \end{aligned}$$

We can determine their coefficients by applying the condition that the total amplitude be invariant under gauge transformations of the photon field. This implies that the expression must vanish when any of the two photons are longitudinal: $A_\mu = \partial_\mu \Lambda$, which means

$$(169) \quad p_\alpha \Gamma_\mu^{\alpha,\beta}(p, q) = 0; \quad q_\beta \Gamma_\mu^{\alpha,\beta}(p, q) = 0.$$

Since

$$(170) \quad \begin{aligned} p_\alpha \delta_1 \Gamma_\mu^{\alpha,\beta}(p, q) &= 0; & q_\beta \delta_1 \Gamma_\mu^{\alpha,\beta}(p, q) &= A_{\mu,\alpha}; \\ p_\alpha \delta_2 \Gamma_\mu^{\alpha,\beta}(p, q) &= A_{\mu,\beta}; & q_\beta \delta_2 \Gamma_\mu^{\alpha,\beta}(p, q) &= 0; \\ A_{\mu,\alpha} &= \varepsilon_{\mu\alpha\beta\gamma} p_\gamma q_\beta, \end{aligned}$$

this fixes the coefficients in front of $\delta_1 \Gamma$ and $\delta_2 \Gamma$.

When now we investigate whether this amplitude is also transversal with respect to the axial vector current, we are struck by a surprise. The counter terms, fixed by condition (170), also contribute here:

$$(171) \quad k_\mu \delta_1 \Gamma_\mu^{\alpha,\beta}(p, q) = -k_\mu \delta_2 \Gamma_\mu^{\alpha,\beta}(p, q) = A_{\alpha,\beta},$$

but they do not cancel against the contribution of the finite part. After imposing gauge-invariance with respect to the two vector insertions, one finds (in the case

of a single chiral fermion)

$$(172) \quad k_\mu \tilde{\Gamma}_\mu^{\alpha,\beta}(p, q) = (4\pi^2)^{-1} \varepsilon_{\mu\alpha\beta\gamma} p_\mu q_\gamma,$$

and this can be rewritten as an effective divergence property of a vector current:

$$(173) \quad \partial_\mu J_\mu^5 = -\frac{iLe^2}{8\pi^2} F_{\mu\nu} \tilde{F}_{\mu\nu},$$

where $\tilde{F}_{\mu\nu} = \frac{1}{2} \varepsilon_{\mu\nu\alpha\beta} F_{\alpha\beta}$, L is the number of left-handed minus the number of right-handed fermions, and it was assumed that the photons couple with charges e .

What is surprising about this is, that the triangle diagram itself, Fig. 5a, appears to be totally symmetric under all permutations, since γ^5 can be permuted to any of the other end-points. Imposing gauge-invariance at two of its three end-points implies breaking of the invariance at the third.

This result is very important. It implies an *induced violation of a conservation law*, apparently to be attributed to the regularization procedure. It also means that it is not possible to couple three *gauge bosons* to such a triangle graph, because this cannot be done in a gauge-invariant way. In most theories, however, we have couplings both to left-handed and to right-handed fermions. Their contributions are of opposite sign, which means that they can cancel out. Therefore, one derives an important constraint on gauge theories with chiral fermions: *The triangle anomalies must cancel out.*

Let the left handed chiral fermions be in representations of the total set of gauge groups that transform as

$$(174) \quad \psi_L^i \rightarrow \psi_L^i + i\Lambda^a T_L^a{}^i{}_j \psi_L^j,$$

where Λ is infinitesimal, and T_L^a are the gauge generators for the left-handed fermions. Similarly for the right-handed ψ_R^i . Define

$$(175) \quad d_L^{abc} = \text{Tr}(T_L^a T_L^b T_L^c + T_L^b T_L^a T_L^c),$$

and similarly d_R^{abc} . The anomaly constraint is then

$$(176) \quad \sum d_L^{abc} = \sum d_R^{abc},$$

where the sum is over all fermion species in the theory. In the Standard Model, the only contributions could come if either one or all three indices of d^{abc} refer to the $U(1)$ group. One quickly verifies that indeed the $U(1)$ charges of the quarks and leptons are distributed in such a way that (176) is completely verified, but *only if*

the number of quark generations and lepton generations are equal. In Subsection 10.3, we will see the physical significance of this observation.

Note, that in the non-Abelian case, there are also anomalies in diagrams with 4 external legs, see Fig. 5(b). They arise from the trilinear terms in $F_{\mu\nu}\tilde{F}_{\mu\nu}$ (the quadrilinear terms cancel). These are the *only* cases where the regularization procedure may violate gauge invariance. In diagrams with more loops, or sub diagrams with more external lines, regularization procedures could be found that preserve gauge invariance.

9 ASYMPTOTIC FREEDOM

9.1 The Renormalization group

It was observed by Stueckelberg and Peterman [Stueckelberg and Peterman, 1953] in 1953, that, although the perturbative expansion of a theory depends on how one splits up the bare parameters in the Lagrangian into lowest order parameters, and counter terms required for the renormalization, the entire theory should not depend on this. This they interpreted as an invariance, and the action of replacing parameters from lowest order to higher order corrections as a group operation. One obtains the ‘Renormalization Group’.

There is only one instance where such transformations really matter, and that is when one compares a theory at one mass- or distance-scale to the same theory at a different scale. A scale transformation must be associated with a replacement of counter terms. Thus, physicists began to identify the notion of a scale transformation as a ‘renormalization group transformation’.

Gell-Mann and Low [Gell-Mann and Low, 1954] observed that this procedure can be used to derive the small-distance behavior of QED. One finds that the effective fine-structure constant depends on the scale μ , described by the equation

$$(177) \quad \frac{\mu^2 d}{d\mu^2} \alpha(\mu) = \beta(\alpha) = \frac{\alpha^2 N}{3\pi} + \mathcal{O}(\alpha^3),$$

where N is the number of charged fermion types. As long as $\alpha(\mu)$ stays small, so that the $\mathcal{O}(\alpha^3)$ terms can be neglected, we see that its μ -dependence is

$$(178) \quad \alpha(\mu) = \frac{\alpha_0}{1 - (\alpha_0 N / 3\pi) \log(\mu^2 / \mu_0^2)}, \quad \text{if } \alpha(\mu_0) = \alpha_0.$$

Things run out of control when μ reaches values comparable to $\exp(3\pi/2N\alpha_0)$, but, at least in the case of QED, where $\alpha_0 \approx 1/137$, this mass scale is so large that in practice no problems are expected. The pole in Eq. (178) is called the Landau pole; Landau concluded that quantum field theories such as QED have no true

continuum limit because of this pole. Gell-Mann and Low suspected, however, that $\beta(\alpha)$ might have a zero at some large value of α , so that, at high values of μ , α approaches this value, but does not exceed this stationary point.

What exactly happens at or near the Landau pole, cannot be established using perturbation expansion alone, since this will depend on *all* higher order terms in Eq. (177); in fact, it is not even known whether Quantum Field Theory can be reformulated accurately enough to decide. The question, however, might be not as important as it seems, since the Landau pole will be way beyond the Planck mass, where we know that gravitational terms will take over; it will be more important to solve Quantum Gravity first.

An entirely different situation emerges in theories where the function $\beta(\lambda)$ is negative. It was long thought that this situation can never arise, unless the coupling strength λ itself is given the wrong sign (the sign that would render the energy density of the classical theory unbounded from below), but this turns out only to be the case in theories that only contain scalar and spinor fields. If there is a *non-Abelian* Yang-Mills component in the theory, negative β functions do occur. In the simplest case, an $SU(2)$ gauge theory with N_f fermions in the elementary doublet representation, the beta function is

$$(179) \quad \frac{\mu^2 d}{d\mu^2} g^2(\mu) = \beta(g^2) = \frac{N_f - 11}{24\pi^2} g^4(\mu) + \mathcal{O}(g^6),$$

so, as long as $N_f < 11$ we have that the coupling strength $g(\mu)$ tends to *zero*, logarithmically, as $\mu \rightarrow \infty$. This feature is called *asymptotic freedom*. In an $SU(N_c)$ gauge theory, the β function is proportional to $N_f - \frac{11}{2}N_c$, so, with the present number of $N_f = 6$ quark flavors, QCD ($N_c = 3$) is asymptotically free. In line with a notation often used, the subscript c here stands for "colour"; in QCD, the number of colours is $N_c = 3$.

9.2 An algebra for the beta functions

In theories with gauge fields, fermions, and scalars, the situation is more complex. A general algorithm for the beta functions has been worked out. The most compact notation for the general result can be given by writing how the entire Lagrange density \mathcal{L} scales under a scale transformation. Let the Lagrangian \mathcal{L} be

$$(180) \quad \begin{aligned} \mathcal{L} = & -\frac{1}{4}G_{\mu\nu}^a G_{\mu\nu}^a - \frac{1}{2}(D_\mu\phi_i)^2 - V(\phi) - \bar{\psi}_i\gamma_\mu D_\mu\psi_i \\ & - \bar{\psi}_i\left(S_{ij}(\phi) + i\gamma_5 P_{ij}(\phi)\right)\psi_j, \end{aligned}$$

where the covariant derivatives are defined as follows:⁸

$$(181) \quad D_\mu\phi_i \equiv \partial_\mu\phi_i + iT_{ij}^a A_\mu^a\phi_j; \quad D_\mu\psi_i = \partial_\mu\psi_i + iU_{ij}^a A_\mu^a\psi_j,$$

⁸ T and U are hermitean, but since ϕ is real, the elements of T must be imaginary.

and the structure constants f^{abc} are defined by

$$(182) [T^a, T^b] = -if^{abc}T_c,$$

so that

$$(183) G_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + f^{abc}A_\mu^b A_\nu^c.$$

We split the fermions into right- and left-handed representations, so that

$$(184) U^a = U_L^a P_L + U_R^a P_R; \quad P_L = \frac{1 + \gamma_5}{2}, \quad P_R = \frac{1 - \gamma_5}{2}.$$

The functions $S(\phi)$ and $P(\phi)$ are at most linear in ϕ and $V(\phi)$ is at most quartic. The Lagrangian (180) is the gauge-invariant part; we do not write the gauge-fixing part or the ghost; the final result will not depend on those details.

The result of an algebraical calculation is that

$$16\pi^2 \frac{\mu^2 d\mathcal{L}}{d\mu^2} = G_{\mu\nu}^a G_{\mu\nu}^b \left[-\frac{11}{12} C_1^{ab} + \frac{1}{24} C_2^{ab} + \frac{1}{6} C_3^{ab} \right] - \Delta V - \bar{\psi}(\Delta S + i\gamma_5 \Delta P)\psi, \quad (185)$$

in which

$$(186) C_1^{ab} = f^{apq} f^{bpq},$$

$$(187) C_2^{ab} = \text{Tr}(T^a T^b),$$

$$(188) C_3^{ab} = \text{Tr}(U_L^a U_L^b + U_R^a U_R^b),$$

$$\Delta V = \frac{1}{4} V_{ij}^2 - \frac{3}{2} V_i (T^2 \phi)_i + \frac{3}{4} (\phi T^a T^b \phi)^2$$

$$(189) \quad + \phi_i V_j \text{Tr}(S_{,i} S_{,j} + P_{,i} P_{,j}) - \text{Tr}(S^2 + P^2)^2 + \text{Tr}[S, P]^2,$$

where

$$(190) V_i \equiv \frac{\partial V(\phi)}{\partial \phi_i}; \quad S_{,i} \equiv \frac{\partial S}{\partial \phi_i}, \quad \text{etc.},$$

and writing $S + iP = W$, one finds ΔS and ΔP from

$$(191) \quad \begin{aligned} \Delta W &= \frac{1}{4} W_i W_i^* W + \frac{1}{4} W W_i^* W_i + W_i W^* W_i \\ &- \frac{3}{2} (U_R^2 W) - \frac{3}{2} W (U_L)^2 + W_i \phi_j \text{Tr}(S_i S_j + P_i P_j). \end{aligned}$$

This expression does not include information on how fields ϕ_i and ψ_i transform under scaling. The fields are not directly observable.

This algebraic expression can be used to find how, in general, coupling strengths run under rescalings of the momenta. It is an interesting exercise to work out what the conditions are for asymptotic freedom, that is, for all coupling strengths to run to zero at infinite momentum. In general, one finds that scalar fields can only exist if there are also gauge fields and fermions present; the latter must be in sufficiently high representations of the gauge group.

10 TOPOLOGICAL TWISTS

The Lagrangian (180) is the most general one allowed if we wish to limit ourselves to coupling strengths that run *logarithmically* under rescalings of the momenta, see for instance Eq. (178). Such theories have a domain of validity that range over exponentially large values of the momenta (in principle over all momenta if the theory is asymptotically free). The most striking feature of this general Lagrangian is that it is topologically highly non-trivial. Locally stable field configurations may exist that have some topological twist in them. In particular, this can be made explicit in the case of a Brout-Englert-Higgs mechanism. Here, these twists can already be seen at the classical level (i.e., ignoring quantum effects).

If we say that a scalar field ϕ_i has a vacuum expectation value, then this means that we perform our perturbation expansion starting with a field value of the form $\phi_i = (F, 0, \dots)$ in the vacuum, after which field fluctuations $\delta\phi$ around this value are assumed to be small. One assumes that the potential $V(\phi)$ has its minimum there. This may appear to violate gauge-invariance, if ϕ_i transform into each other under local gauge transformations, but strictly speaking the phrase “spontaneous breakdown of local gauge symmetry” is inappropriate, because it may also simply mean that we choose a gauge condition. It is however a fact that the spectrum of physical particles comes out to be altogether different if we perturb around $\phi_i = 0$, so this ‘Higgs mode’ is an important notion in any case.

10.1 Vortices

If the Higgs field has only two real components (such as when $U(1)$ is broken into the identity group), one may consider a configuration where this field makes a full twist over 360° when following a closed curve. Inside the curve there must be a zero. The zeros must form a curve themselves, and they cost energy. This is the Abrikosov vortex. Away from its center, one may transform ϕ_i back to a constant value, but this generates a vector potential $A_\mu(x)$, obeying

$$(192) \quad \oint A_\mu dx^\mu = \frac{2\pi}{e},$$

which means that this vortex carries an amount of magnetic flux, of magnitude exactly $2\pi/e$. Apparently, in this model, magnetic field lines condense into locally stable vortices [Nielsen and Olesen, 1973]. This is also what happens to magnetic fields inside a superconductor.

10.2 Magnetic Monopoles

Something similar may happen if the Higgs field has three real components. In that case, one can map the S_2 sphere of minima of $V(\phi)$, onto a sphere in 3-space. There will be isolated zeros inside this sphere. These objects behave as locally stable particles. If one tries to transform the field locally to a constant value, one finds that a vector potential again may emerge.

If, for example, in an $SU(2)$ theory, a Higgs in the adjoint representation (which has 3 real components) breaks the gauge group down to $U(1)$, then one finds the vector potential of an isolated magnetic source inside the sphere. This means that the source is a magnetic monopole with magnetic charge $g_m = \frac{4\pi}{e}$, where e is the original coupling strength of the $SU(2)$ theory. Indeed, Dirac [1931; 1948] has derived, back in 1931, that magnetic charges g_m and electric charges q must obey the *Dirac quantization condition*

$$(193) \quad q g_m = 2\pi n.$$

Apparently, for the monopole in this model, $n = 2$. However, it is easy to introduce particles in the elementary representation, which have $q = \frac{1}{2}e$; these then saturate the Dirac condition (193).

Dirac could not say much about the mass of his magnetic monopoles. In the present theories, however, the mass is calculable. In general, the magnetic monopole mass turns out to be the mass of an ordinary particle divided by a number of the order of the gauge coupling strength squared.

Careful analysis of the existing Lie groups and the way they may be broken spontaneously into one or more subgroups $U(1)$, reveals a general feature: *Only if the underlying gauge group is compact, and has a compact covering group, must electric charges in the $U(1)$ gauge groups be quantized* (otherwise, it would not be forbidden to add arbitrary real numbers to the $U(1)$ charges), *and whenever the covering group of the underlying gauge group is compact, magnetic monopole solutions can be constructed*. Apparently, whenever the gauge group structure provides for a compelling reason for electric charges to be quantized, the existence of magnetic monopole solutions is guaranteed. Thus, assuming that Nature has compelling reasons for the charge units of electrons and protons to be equal, and quantized into multiples of e , we must assume that magnetic monopole solutions must exist. However, in most 'Grand Unification Schemes', the relevant mass scale is many orders of magnitude higher than the mass scale of particles studied today,

so the monopoles, whose mass is that divided by a coupling strength squared, are even heavier.

From the structure of the Higgs field of a monopole, one derives that the system is invariant under rotations provided that rotations are associated with gauge rotations. A consequence of this is, that elementary particles with half-odd isospin, when bound to a monopole, produce bound states with half-odd integer *orbital* angular momentum [Hasenfratz and 't Hooft, 1976]. What is strange about this, is that such particles should develop Dirac statistics. Indeed, one can derive that both the spin and the statistics of bound states of electric and magnetic charges, flip from Bose-Einstein to Fermi-Dirac or back [Goldhaber, 1976] if they form odd values of the Dirac quantum n (Eq. 193).

10.3 Instantons

A Higgs field with two real components gives rise to vortices, a Higgs with three components gives magnetic monopoles, so what do we get if a Higgs field has four real components? This is the case if, for instance, $SU(2)$ is broken spontaneously into the identity by a Higgs in the fundamental representation (two complex = 4 real components). The topologically stable objects one finds are stable points in four-dimensional space-time. They represent *events*, and, referring to their particle-like appearance, the resulting solutions (in Euclidean space) were called ‘instantons’. Because this Higgs field, in the case of $SU(2)$, breaks the gauge symmetry completely, one can argue that this solution is also topologically stable in pure gauge theories, without a Higgs mechanism at all. Far from the origin, the vector potential field is described as a local gauge rotation of the value $A_\mu^a(x) = 0$. The gauge rotation in question, $\Omega(x)$, is described by noting that the $SU(2)$ matrices form an S_3 space, *i.e.*, the three dimensional surface of a sphere in four dimensions. Mapping this S_3 one-to-one onto the boundary of some large region in (Euclidean) space-time, gives the field configuration of an instanton.

It was noted by Belavin, Polyakov, Schwarz and Tyupkin [Belavin et al., 1975] (who also were the first to write down this solution) that this solution has a non-vanishing value of

$$(194) \int d^4x F_{\mu\nu}^a \tilde{F}_{\mu\nu}^a = \frac{32\pi^2}{g^2}.$$

The integrand is the divergence of a current:

$$(195) F_{\mu\nu}^a \tilde{F}_{\mu\nu}^a = \partial_\mu K_\mu; \quad K_\mu = 2\varepsilon_{\mu\nu\alpha\beta} A_\nu^a (\partial_\alpha A_\beta^a + \frac{1}{3}gf^{abc} A_\alpha^b A_\beta^c),$$

the so-called Chern-Simons current. This current, however, is not gauge-invariant, which is why it does not vanish at infinity. It does vanish after the gauge transformation $\Omega(x)$ that replaces A_μ^a at infinity by 0. Eq. (194) is most easily derived

by using this Chern-Simons current. It so happens that the instanton is also a solution of the equation

$$(196) F_{\mu\nu} = \tilde{F}_{\mu\nu},$$

so that we also find the action to be given by $-8\pi^2/g^2$.

In a pure gauge theory (one without fermions), instantons can be interpreted as representing *tunneling transitions*. In ordinary Quantum Mechanics, tunneling is an exponentially suppressed transition. The exponential suppression is turned into an oscillating expression if we replace time t by an imaginary quantity $i\tau$. The oscillating exponent is the action of a classical transition in imaginary time. One may also say that a tunneling transition can be described by a classical mechanical transition if the potential $V(\vec{q})$ is replaced by $2E - V(\vec{q})$, where E is the energy. The classical action then represents the quantity in the exponent of the (exponentially suppressed) tunneling transition.

The above substitution is exactly what one gets by replacing time t by $i\tau$. In relativistic Quantum Field Theory, this is also exactly the Wick rotation from Minkowski space-time into Euclidean space-time. In short, instantons represent tunneling that is associated with the suppression factor $e^{-8\pi^2/g^2}$.

The transition can be further understood by formulating a gauge theory in the temporal gauge, $A_0 = 0$. In this gauge, there is a residual invariance under gauge transformations $\Lambda(\mathbf{x})$ that are time-independent. All 'physical states', therefore, come as representations of this local gauge group. Normally, however, we restrict ourselves to the trivial representation, $\Omega|\psi\rangle = |\psi\rangle$, where $\Omega = e^{i\int \Lambda(\mathbf{x})d^3\mathbf{x}}$, because this configuration is conserved in time, and because any other choice would violate Lorentz invariance. However, closer analysis shows that one only has to impose this constraint for those gauge transformations that can be continuously reached from the identity transformation. This is *not* the case for transformations obtained by mapping the S_3 space of the $SU(2)$ transformations non-trivially onto three-space \mathbb{R}_3 . These transformations form a discrete set, characterized by the integers $k = 0, \pm 1, \pm 2, \dots$. Writing

$$(197) \Omega_k(\mathbf{x}) = \Omega_1(\mathbf{x})^k, \quad \Omega_k|\psi\rangle = e^{i\theta k}|\psi\rangle,$$

we find that the tunneling transitions described by instantons cause an exponentially suppressed θ dependence of physical phenomena in the theory. Since, under parity transformations P , the angle θ turns into $-\theta$, a non-vanishing θ also implies an explicit parity (eventually, PC) violation of the strong interactions.

In the presence of fermions, the situation is altogether different. Due to the chiral anomaly, we have for the current of chiral fermions $J_\mu^5(x)$, the equation (173). The total number of chiral fermions, $Q^5 = \int d^3\mathbf{x} J_0^5(\mathbf{x})$ changes by one unit due to an instanton: $\Delta Q^5 = \pm 1$. This can be understood by noting that the Dirac

equation for massless, chiral fermions has one localized solution in the Euclidean space of an instanton. In Minkowski space-time, this solution turns into a state that describes a chiral fermion either disappearing into the Dirac sea, or emerging from it, so that, indeed, the number of particles minus anti-particles changes by one unit for every chiral fermion species. If left- and right handed fermions are coupled the same way to the gauge field, as in QCD, the instanton removes a left-handed fermion and creates a right-handed one, or, in other words, it flips the chirality. This $\Delta Q^5 = \pm 2$ event has exactly the quantum numbers of a mass term for the Goldstone boson that would be associated to the conservation of chiral charge, the η particle. This explains why the η particle is considerably heavier than the pions, which have lost most of their mass due to chiral symmetry of the quarks [’t Hooft, 1986].

What one concludes from the study of instantons is that QCD, the theory for the strong interactions, neatly explains the observed symmetry structure of the hadron spectrum, including the violation of chiral charge conservation that accounts for the η mass.

In the electro-weak sector, one also has instantons. We now see that the cancellation of the anomalies in the quark and the lepton sector implies an important property of the electro-weak theory: since the anomalies do not respect gauge-invariance of the quark sector alone, quarks can be shown not to be exactly conserved. One finds that instantons induce baryon number violating events: three baryons (nine different quarks all together) may transmute into three anti-leptons, or vice versa.

11 CONFINEMENT

An important element in the Standard Model is the gauge theory for the strong interactions, based on the gauge group $SU(3)$. Quarks are fermions in the elementary representation of $SU(3)$. The observed hadronic particles all are bound states of quarks and/or anti-quarks, in combinations that are gauge-invariant under $SU(3)$. An important question is: what is the nature of the forces that binds these quarks together? We have seen that vortex solutions can be written down that would cause an interesting force pattern among magnetic monopoles: in a Higgs theory with magnetic monopoles, these monopoles could be bound together with Abrikosov vortices.

Indeed, this would be a confining force: every magnetic monopole must be the end point of a vortex, whose other end point is a monopole of opposite magnetic charge. Indeed, the confinement would be *absolute*: isolated monopoles cannot exist. It was once thought that, therefore, quarks must be magnetic monopoles. This, however, would be incompatible with the finding that quarks only interact weakly at small distances, magnetic charges being always quite strong. A more elegant idea is that the binding force forms electric rather than magnetic vortices.

An electric vortex can be understood as the *dual transformation* of a magnetic vortex. It comes about when the Brout-Englert-Higgs mechanism affects freely moving magnetically charged particles. Further analytic arguments, as well as numerical investigations, have revealed that indeed such objects are present in QCD, and that the Higgs mechanism may occur in this sector. Let us briefly explain the situation in words.

11.1 *The maximally Abelian gauge*

A feature that distinguishes non-Abelian gauge theories from Abelian ones, is that a reference frame for the gauge choice, the gauge condition, can partly be fixed locally in terms of the pure gauge fields alone; noticing that the covariant field strengths $G_{\mu\nu}$ transform as the adjoint representation, one may choose the gauge such that one of these components, say G_{12} , is diagonal. This then removes the non-Abelian part of the gauge group, but the diagonal part, called the Cartan subgroup, remains. In this way, a non-Abelian gauge theory turns into an Abelian one. A slightly smarter, but non-local gauge that does the same is the condition that $\sum_{i \neq j} (A_{\mu j}^i)^2$ is minimized. It is called the maximally Abelian gauge.

However, such a gauge choice does produce singularities. These typically occur when two eigenvalues of G_{12} coincide. It is not difficult to convince oneself that these singularities behave as particles, and that these particles carry magnetic charges with respect to the Cartan subgroup. *Absolute confinement occurs as soon as these magnetically charged particles undergo a Brout-Englert-Higgs mechanism.*

Although this still is the preferred picture explaining the absolute nature of the quark confining force, it may be noted that the magnetically charged particles do not have to be directly involved with the confinement *mechanism*. Rather, they are indicators. This, we deduce from the fact that confinement also occurs in theories with a very large number N_c of colors; in the limit $N_c \rightarrow \infty$, magnetically charged particles appear to be suppressed in the perturbative regime, but the electric vortices are nevertheless stable. The strength of a vortex is determined by its finite width, and this width is controlled by the lightest gluonic state, the 'glueball'. At distance scales large compared to the inverse mass of the lightest glueball, an electric vortex cannot break.

Confinement is a condensation phase that is a logical alternative of the Brout-Englert-Higgs phase. In some cases, however, these two phases may coexist. An example of such a coexistence is the $SU(2)$ sector of the Standard Model. Conventionally, this sector is viewed as a prototype of the Higgs mechanism, but it so happens that the $SU(2)$ sector of the Standard Model can be treated exactly like the colour $SU(3)$ sector: as if there is confinement. To see this, one must observe that the Higgs doublet field can be used to fix the $SU(2)$ sector of the gauge group unambiguously. This means that all physical particles can be connected to gauge-invariant sources by viewing them as gauge-invariant bound states of the Higgs

particle with the other elementary doublets of the model. For instance, writing the Higgs doublet as $\phi_a = \begin{pmatrix} F \\ 0 \end{pmatrix} + \tilde{\phi}_a$, and the lepton doublet as ψ_a , the electron is seen to be associated to the ‘baryonic’ field $\varepsilon^{ab}\phi_a\psi_b$, the neutrino is $\phi^{*a}\psi_a$, the Z_0 boson is $\phi^{*a}D_\mu\phi_a$, and so on.

Theories in which the confinement phase is truly distinct from the Higgs phase are those where the Higgs field is not a one-to-one representation of the gauge group, such as the adjoint representation of $SU(2)$.

12 OUTLOOK

Quantum Field Theory has reached a respectable status as an accurate and well-studied description of sub-atomic particles. From a purely mathematical point of view, there are some inherent limitations to the accuracy by which it defines the desired amplitudes, but in nearly all conceivable circumstances, its intrinsic accuracy is much higher than what can be reached in experiments. This does not mean that we can reach such accuracy in real calculations, which more often than not suffer from technical limitations, particularly where the interactions are strong, as in QCD. In this domain, there is still a need for considerable technical advances.

12.1 *Naturalness*

When the Standard Model, as known today, is extrapolated to energy domains beyond approximately 1 TeV, a difficulty is encountered that is not of a mathematical nature, but rather a physical one: it becomes difficult to believe that it represents the real world. The bare Lagrangian, when considered on a very fine lattice, is required to have parameters that must be tuned very precisely in order to produce particles such as the Higgs particle and the weak vector bosons, whose masses are much *less* than 1 TeV. This fine-tuning is considered to be *unnatural*. In a respectable physical theory, such a coincidence is not expected. With some certainty, one can state that the fundamental laws of Nature must allow for a more elegant description at high energies than a lattice with such fine-tuning. What is generally expected is either a new symmetry principle or possibly a new regime with an altogether different set of physical fields.

A candidate for a radically different regime is the so-called *technicolour* theory, a repetition of QCD but with a typical energy scale of a TeV rather than a GeV. The quarks, leptons and Higgs particles of the Standard Model would then all turn out to be the hadrons of this technicolour theory. Different gauge groups could replace $SU(3)$ here. However, according to this scheme, a new strong interaction regime would be reached, where perturbation expansions used in the weak sector of the Standard Model would have to break down. As precision measurements and

calculations continue to confirm the reliability of these perturbation expansions, the technicolour scenario is considered to be unlikely.

12.2 *Supersymmetry*

A preferred scenario is a simple but beautiful enhancement of the symmetries of the Standard Model: supersymmetry. This symmetry, which puts fermions and bosons into single multiplets, does not really modify the fundamental aspects of the theory. But it does bring about considerable simplifications in the expressions for the amplitudes, not only in the perturbative sector, but also, in many cases, it allows us to look deeper into the non-perturbative domains of the theories. There is a vast amount of literature on supersymmetry, but some aspects of it are still somewhat obscure. We would like to know more about the physical origin and meaning of supersymmetry, as well as the mechanism(s) causing it to be broken — and made almost invisible — at the domain of the Standard Model that is today accessible to experimental observation.

12.3 *Resummation of the Perturbation Expansion*

The perturbation expansion in Quantum Field Theory is almost certain to be divergent for any value of the coupling parameter(s). A simple argument for its divergence has been put forward by Dyson [Dyson, 1952]: imagine that in the theory of QED there were a bound ε such that, whenever $|\alpha| < \varepsilon$, where α is the fine-structure constant, perturbation expansions would converge. Then it would converge for some negative real value of α . However, one can easily ascertain that for any negative value of α , the vacuum would be unstable: vacuum fluctuations would allow large numbers of electrons to be pair-created, and since like charges attract, highly charged clouds of electrons could have negative energies.

Theories with asymptotic freedom may allow for a natural way to re-sum the perturbation series, by first solving the theory at high energy with extreme precision, after which one has to integrate the Schrödinger equation to obtain the physical amplitudes at lower energy. Such a program has *not* yet been carried out, because integrating these Schrödinger equations is beyond our present capabilities, but one may suspect that, as a matter of principle, it should be possible. Theories that are not asymptotically free may perhaps allow for more precise treatments if an ultra-violet fixed point can be established.

The extent of the divergence of the perturbation expansion can be studied or predicted. This one does using the Borel resummation technique. An amplitude

$$(198) \Gamma(\lambda) = \sum_{n=1}^{\infty} a_n \lambda^n ,$$

can be rewritten as

$$\Gamma(\lambda) = \int_0^\infty B(z)e^{-z/\lambda} dz,$$

$$(199) B(z) = \sum_{k=0}^\infty a_{k+1} z^k / k! .$$

The series for $B(z)$ is generally expected to have a finite radius of convergence. If $B(z)$ can be analytically extended to the domain $0 \leq z < \infty$, then that (re-) defines our amplitude. In general, however, one can derive that $B(z)$ must have singularities on the real axis, for instance where z corresponds to the action of instantons or instanton pairs. In addition, singularities associated to the infrared and/or ultraviolet divergences of the theory are expected. Sometimes, these different singularities interfere.

12.4 General Relativity and Superstring Theory

It is dubious, however, whether the issue of convergence or divergence of the perturbation expansion is of physical relevance. We *know* that Quantum Field Theory cannot contain the entire truth concerning the sub-atomic world; the gravitational force is guaranteed not to be renormalizable, so at those scales where this force becomes comparable to the other forces, the so-called Planck scale, a radically new theory is called for. *Superstring Theory* is presently holding the best promise to evolve into such a theory. With this theory, physicists are opening a new chapter, where we leave conventional Quantum Field Theory, as described in this paper, behind. In its present form, Superstring Theory appears to have turned into a collection of wild ideas called *M*-theory, whose foundations are still extremely shaky. Some of the best minds of the world are competing to turn this theory into something that can be used to provide for reliable predictions and that can be taught in a text book, but this has not yet been achieved.

BIBLIOGRAPHY

- [Adler, 1969] S. Adler. Axial-vector vertex in spinor electrodynamics. *Physical Review*, 177: 2426–2438, 1969.
- [Adler and Bardeen, 1969] S. Adler and W. Bardeen. Absence of higher-order corrections in the anomalous axial-vector divergence equation. *Physical Review*, 182: 1517–1536, 1969.
- [Aitchison and Hey, 1989] I. Aitchison and A. Hey. *Gauge Theories in Particle Physics*. Adam Hilger, 1989.
- [Bardeen, 1969] W. Bardeen. Anomalous Ward identities in spinor field theories. *Physical Review*, 184: 1848–1859, 1969.
- [Becchi *et al.*, 1975] C. Becchi, A. Rouet, and R. Stora. Renormalization of the Abelian Higgs-Kibble model. *Communications in Mathematical Physics*, 42: 127–162, 1975.
- [Becchi *et al.*, 1976] C. Becchi, A. Rouet, and R. Stora. Renormalization of gauge theories. *Annals of Physics*, 98: 287–321, 1976.

- [Belavin et al., 1975] A. Belavin, A. Polyakov, A. Schwartz, and Y. Tyupkin. Pseudoparticle solutions of the Yang-Mills equations. *Physics Letters*, 59 B: 85–87, 1975.
- [Bell and Jackiw, 1969] J. Bell and R. Jackiw. A pcac puzzle: π /sup 0/ to $\gamma\gamma$ in the sigma-model. *Il Nuovo Cimento*, 60A: 47–61, v.
- [Cheng and Li, 1984] T. Cheng and L. Li. *Gauge theory of elementary particle physics*. Clarendon Press, Oxford, 1984.
- [Crease and Mann, 1986] R. Crease and C. Mann. *The Second Creation: makers of the revolution in twentieth-century physics*. Macmillan, New York, 1986.
- [de Wit and Smith, 1986] B. de Wit and J. Smith. *Field Theory in Particle Physics*. North Holland, 1986.
- [DeWitt, 1964] B. DeWitt. Theory of radiative corrections for non-Abelian gauge fields. *Physical Review Letters*, 12: 742–746, 1964.
- [DeWitt, 1967a] B. DeWitt. Quantum theory of gravity. I. The canonical theory. *Physical Review*, 160: 1113–1148, 1967a.
- [DeWitt, 1967b] B. DeWitt. Quantum theory of gravity. II. The manifestly covariant theory. *Physical Review*, 162: 1195–1239, 1967b.
- [Dirac, 1931] P. Dirac. Quantised singularities in the electromagnetic field. *Proceedings of the Royal Society*, A133: 60–72, 1931.
- [Dirac, 1948] P. Dirac. The theory of magnetic poles. *Physical Review*, 74: 817–830, 1948.
- [Dyson, 1952] F. Dyson. Divergence of perturbation theory in quantum electrodynamics. *Physical Review*, 85: 631–632, 1952.
- [Englert and Brout, 1964] F. Englert and R. Brout. Broken symmetry and the mass of gauge vector mesons. *Physical Review Letters*, 13: 321–323, 1964.
- [Faddeev, 1969] L. Faddeev. The Feynman integral for singular Lagrangians. *Theoretical and Mathematical Physics*, 1:3–18, 1969. Russian. English summary.
- [Faddeev and Popov, 1967] L. Faddeev and V. Popov. Feynman diagrams for the Yang-Mills field. *Physics Letters*, 25B: 27–29, 1967.
- [Faddeev and Takhtajan, 1984] L. Faddeev and L. Takhtajan. *Soviet Journal of Mathematics*, 24: 241, 1984.
- [Gell-Mann and Low, 1954] M. Gell-Mann and F. Low. Quantum electrodynamics at small distances. *Physical Review*, 95: 1300–1312, 1954.
- [Goldhaber, 1976] A. Goldhaber. Connection of spin and statistics for charge-monopole composites. *Physical Review Letters*, 36: 1122–1125, 1976.
- [Hasenfratz and 't Hooft, 1976] P. Hasenfratz and G. 't Hooft. Fermion-boson puzzle in a gauge theory. *Physical Review Letters*, 36: 1119–1122, 1976.
- [Higgs, 1964a] P. Higgs. Broken symmetries and the masses of gauge bosons. *Physical Review Letters*, 13: 508–509, 1964a.
- [Higgs, 1964b] P. Higgs. Broken symmetries, massless particles and gauge fields. *Physics Letters*, 12: 132–133, 1964b.
- [Higgs, 1966] P. Higgs. Spontaneous symmetry breakdown without massless bosons. *Physical Review*, 145: 1156–1163, 1966.
- [Hoddeson et al., 1997] L. Hoddeson, L. Brown, M. Riordan, and M. Dresden (eds.). *The Rise of the Standard Model, Particle Physics in the 1960s and 1970s*. Cambridge University Press, 1997.
- [Itzykson and Zuber, 1980] C. Itzykson and B. Zuber. *Quantum Field Theory*. McGraw Hill, New York, 1980.
- [Jackiw, 1995] R. Jackiw. *Diverse topics in Theoretical and Mathematical Physics*. World Scientific, 1995.
- [Nielsen and Olesen, 1973] H. Nielsen and P. Olesen. Vortex-line models for dual strings. *Nuclear Physics B*, 61: 45–61, 1973.
- [Pais, 1986] A. Pais. *Inward bound: of matter and forces in the physical world*. Oxford University Press, 1986.
- [Ryder, 1985] L. Ryder. *Quantum Field Theory*. Cambridge Univ. Press, 1985.
- [Slavnov, 1972] A. Slavnov. Ward identities in gauge theories. *Theoretical and Mathematical Physics*, 10: 99–107, 1972.
- [Stueckelberg and Peterman, 1953] E. Stueckelberg and A. Peterman. *Helvetica Physica Acta*, 26: 499–, 1953.
- ['t Hooft, 1971] G. 't Hooft. Renormalization of massless Yang-Mills fields. *Nuclear Physics B*, 33: 173–199, 1971.

- [’t Hooft, 1986] G. ’t Hooft. How instantons solve the $U(1)$ problem. *Physics Reports*, 142: 357–387, 1986.
- [’t Hooft and Veltman, 1972a] G. ’t Hooft and M. Veltman.. Combinatorics of gauge fields. *Nuclear Physics B*, 50: 318–353, 1972a.
- [’t Hooft and Veltman, 1972b] G. ’t Hooft and M. Veltman. Regularization and renormalization of gauge fields. *Nuclear Physics B*, 44: 189–213, 1972b.
- [’t Hooft and Veltman, 1994] G. ’t Hooft and M. Veltman. *Under the Spell of the Gauge Principle*, volume 19 of *Advanced Series in Mathematical Physics*. World Scientific, Singapore, 1994.
- [Taylor, 1971] J. Taylor. Ward identities and charge renormalization of the Yang-Mills field. *Nuclear Physics B*, 33: 436–444, 1971.
- [Tyutin, 1975] I. Tyutin. Lebedev Preprint. FIAN 39. unpublished, 1975.
- [Yang and Mills, 1954] C. Yang and R. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Physical Review*, 96: 191–195, 1954.

This page intentionally left blank

ALGEBRAIC QUANTUM FIELD THEORY

Hans Halvorson (with an appendix by Michael Müger)

INTRODUCTION

From the title of this Chapter, one might suspect that the subject is some *idiosyncratic* approach to quantum field theory (QFT). The approach is indeed idiosyncratic in the sense of demographics: only a small proportion of those who work on QFT work on *algebraic* QFT (AQFT). However, there are particular reasons why philosophers, and others interested in foundational issues, will want to study the “algebraic” approach.

In philosophy of science in the analytic tradition, studying the foundations of a theory T has been thought to presuppose some minimal level of clarity about the referent of T . (Moreover, to distinguish philosophy from sociology and history, T is not taken to refer to the activities of some group of people.) In the early twentieth century, it was thought that the referent of T must be a set of axioms of some formal, preferably first-order, language. It was quickly realized that not many interesting physical theories can be formalized in this way. But in any case, we are no longer in the grip of axiomania, as Feyerabend called it. So, the standards were loosened somewhat — but only to the extent that the standards were simultaneously loosened within the community of professional mathematicians. There remains an implicit working assumption among many philosophers that studying the foundations of a theory requires that the theory has a mathematical description. (The philosopher’s working assumption is certainly satisfied in the case of statistical mechanics, special and general relativity, and nonrelativistic quantum mechanics.) In any case, whether or not having a mathematical description is mandatory, having such a description greatly facilitates our ability to draw inferences securely and efficiently.

So, philosophers of physics have taken their object of study to be theories, where theories correspond to mathematical objects (perhaps sets of models). But it is not so clear where “quantum field theory” can be located in the mathematical universe. In the absence of some sort of mathematically intelligible description of QFT, the philosopher of physics has two options: either find a new way to understand the task of interpretation, or remain silent about the interpretation of quantum field theory.¹

It is for this reason that AQFT is of particular interest for the foundations of quantum field theory. In short, AQFT is our best story about where QFT lives

¹For the first option, see [Wallace, forthcoming].

in the mathematical universe, and so is a natural starting point for foundational inquiries.

1 ALGEBRAIC PROLEGOMENA

This first section provides a minimal overview of the mathematical prerequisites of the remainder of the Chapter.

1.1 von Neumann algebras

The standard definition of a von Neumann algebra involves reference to a topology, and it is then shown (by von Neumann's double commutant theorem) that this topological condition coincides with an algebraic condition (condition 2 in the Definition 2). But for present purposes, it will suffice to take the algebraic condition as basic.

DEFINITION 1. Let \mathcal{H} be a Hilbert space. Let $\mathfrak{B}(\mathcal{H})$ be the set of bounded linear operators on \mathcal{H} in the sense that for each $A \in \mathfrak{B}(\mathcal{H})$ there is a smallest nonnegative number $\|A\|$ such that $\langle Ax, Ax \rangle^{1/2} \leq \|A\|$ for all unit vectors $x \in \mathcal{H}$. [Subsequently we use $\|\cdot\|$ ambiguously for the norm on \mathcal{H} and the norm on $\mathfrak{B}(\mathcal{H})$.] We use juxtaposition AB to denote the composition of two elements A, B of $\mathfrak{B}(\mathcal{H})$. For each $A \in \mathfrak{B}(\mathcal{H})$ we let A^* denote the unique element of $\mathfrak{B}(\mathcal{H})$ such that $\langle A^*x, y \rangle = \langle x, Ay \rangle$, for all $x, y \in \mathfrak{R}$.

DEFINITION 2. Let \mathfrak{R} be a $*$ -subalgebra of $\mathfrak{B}(\mathcal{H})$, the bounded operators on the Hilbert space \mathcal{H} . Then \mathfrak{R} is a von Neumann algebra if

1. $I \in \mathfrak{R}$,
2. $(\mathfrak{R}')' = \mathfrak{R}$,

where $\mathfrak{R}' = \{B \in \mathfrak{B}(\mathcal{H}) : [B, A] = 0, \forall A \in \mathfrak{R}\}$.

DEFINITION 3. We will need four standard topologies on the set $\mathfrak{B}(\mathcal{H})$ of bounded linear operators on \mathcal{H} . Each of these topologies is defined in terms of a family of seminorms — see [Kadison and Ringrose, 1997, Chaps. 1,5] for more details.

- The uniform topology on $\mathfrak{B}(\mathcal{H})$ is defined in terms of a single norm:

$$\|A\| = \sup\{\|Av\| : v \in \mathcal{H}, \|v\| \leq 1\},$$

where the norm on the right is the given vector norm on \mathcal{H} . Hence, an operator A is a limit point of the sequence $(A_i)_{i \in \mathbb{N}}$ iff $(\|A_i - A\|)_{i \in \mathbb{N}}$ converges to 0.

- The weak topology on $\mathfrak{B}(\mathcal{H})$ is defined in terms of the family $\{p_{u,v} : u, v \in \mathcal{H}\}$ of seminorms where

$$p_{u,v}(A) = \langle u, Av \rangle.$$

The resulting topology is not generally first countable, and so the closure of a subset S of $\mathfrak{B}(\mathcal{H})$ is generally larger than the set of all limit points of sequences in S . Rather, the closure of S is the set of limit points of generalized sequences (nets) in S — see [Kadison and Ringrose, 1997, Chap. 1] for more details. A net $(A_i)_{i \in \mathcal{I}}$ in $\mathfrak{B}(\mathcal{H})$ converges weakly to A just in case $(p_{u,v}(A_i))_{i \in \mathcal{I}}$ converges to $p_{u,v}(A)$ for all $u, v \in \mathcal{H}$.

- The strong topology on $\mathfrak{B}(\mathcal{H})$ is defined in terms of the family $\{p_v : v \in \mathcal{H}\}$ of seminorms where

$$p_v(A) = \|Av\|.$$


Thus, a net $(A_i)_{i \in \mathcal{I}}$ converges strongly to A iff $(p_v(A_i))_{i \in \mathcal{I}}$ converges to $p_v(A)$, for all $v \in \mathcal{H}$.

- The ultraweak topology on $\mathfrak{B}(\mathcal{H})$ is defined in terms of the family $\{p_\rho : \rho \in \mathcal{T}(\mathcal{H})\}$ where $\mathcal{T}(\mathcal{H})$ is the set of positive, trace 1 operators (“density operators”) on \mathcal{H} and

$$p_\rho(A) = \text{Tr}(\rho A).$$

Thus a net $(A_i)_{i \in \mathcal{I}}$ converges ultraweakly to A just in case $(\text{Tr}(\rho A_i))_{i \in \mathcal{I}}$ converges to $\text{Tr}(\rho A)$, for all $\rho \in \mathcal{T}(\mathcal{H})$.

FACT 4. The topologies are ordered as follows:



Since closed sets are just the complements of open sets, this means that a weakly closed set is ultraweakly closed, and an ultraweakly closed subset is norm closed. Furthermore, the four topologies on $\mathfrak{B}(\mathcal{H})$ coincide iff \mathcal{H} is finite dimensional.

FACT 5. If S is a bounded, convex subset of $\mathfrak{B}(\mathcal{H})$, then the weak, ultraweak, and strong closures of S are the same.

FACT 6. For a $*$ -algebra \mathfrak{A} on \mathcal{H} that contains I , the following are equivalent: (i) \mathfrak{A} is weakly closed; (ii) $\mathfrak{A}'' = \mathfrak{A}$. This is von Neumann’s double commutant theorem.

DEFINITION 7. Let \mathfrak{A} be a subset of $\mathfrak{B}(\mathcal{H})$. A vector $x \in \mathcal{H}$ is said to be *cyclic* for \mathfrak{A} just in case $[\mathfrak{A}x] = \mathcal{H}$, where $\mathfrak{A}x = \{Ax : A \in \mathfrak{A}\}$, and $[\mathfrak{A}x]$ is the closed linear span of $\mathfrak{A}x$. A vector $x \in \mathcal{H}$ is said to be *separating* for \mathfrak{A} just in case $Ax = 0$ and $A \in \mathfrak{A}$ entails $A = 0$.

FACT 8. Let \mathfrak{A} be a von Neumann algebra on \mathcal{H} , and let $x \in \mathcal{H}$. Then x is cyclic for \mathfrak{A} iff x is separating for \mathfrak{A}' .

DEFINITION 9. A *normal* state of a von Neumann algebra \mathfrak{A} is an ultraweakly continuous state. We let \mathfrak{A}_* denote the normal state space of \mathfrak{A} .

1.2 C^* -algebras and their representations

DEFINITION 10. A C^* -algebra is a pair consisting of a $*$ -algebra \mathfrak{A} and a norm $\|\cdot\| : \mathfrak{A} \rightarrow \mathbb{R}$ such that

$$\|AB\| \leq \|A\| \cdot \|B\|, \quad \|A^*A\| = \|A\|^2,$$

for all $A, B \in \mathfrak{A}$. We usually use \mathfrak{A} to denote the algebra and its norm.

In this Chapter, we will only use C^* -algebras that contain a multiplicative identity I .

DEFINITION 11. A state ω on \mathfrak{A} is a linear functional such that $\omega(A^*A) \geq 0$ for all $A \in \mathfrak{A}$, and $\omega(I) = 1$.

DEFINITION 12. A state ω of \mathfrak{A} is said to be *mixed* if $\omega = \frac{1}{2}(\omega_1 + \omega_2)$ with $\omega_1 \neq \omega_2$. Otherwise ω is said to be *pure*.

DEFINITION 13. Let \mathfrak{A} be a C^* -algebra. A *representation* of \mathfrak{A} is a pair (\mathcal{H}, π) , where \mathcal{H} is a Hilbert space and π is a $*$ -homomorphism of \mathfrak{A} into $\mathfrak{B}(\mathcal{H})$. A representation (\mathcal{H}, π) is said to be *irreducible* if $\pi(\mathfrak{A})$ is weakly dense in $\mathfrak{B}(\mathcal{H})$. A representation (\mathcal{H}, π) is said to be *faithful* if π is an isomorphism.

DEFINITION 14. Let (\mathcal{H}, π) and (\mathcal{K}, ϕ) be representations of a C^* -algebra \mathfrak{A} . Then (\mathcal{H}, π) and (\mathcal{K}, ϕ) are said to be:

1. *unitarily equivalent* if there is a unitary $U : \mathcal{H} \rightarrow \mathcal{K}$ such that $U\pi(A) = \phi(A)U$ for all $A \in \mathfrak{A}$.
2. *quasiequivalent* if the von Neumann algebras $\pi(\mathfrak{A})''$ and $\phi(\mathfrak{A})''$ are $*$ -isomorphic.
3. *disjoint* if they are not quasiequivalent.

DEFINITION 15. A representation (\mathcal{K}, ϕ) is said to be a *subrepresentation* of (\mathcal{H}, π) just in case there is an isometry $V : \mathcal{K} \rightarrow \mathcal{H}$ such that $\pi(A)V = V\phi(A)$ for all $A \in \mathfrak{A}$.

FACT 16. Two representations are quasiequivalent iff they have unitarily equivalent subrepresentations.

The famous Gelfand-Naimark-Segal (GNS) theorem shows that every C^* -algebraic state can be represented by a vector in a Hilbert space.

THEOREM 17 (GNS). *Let ω be a state of \mathfrak{A} . Then there is a representation (\mathcal{H}, π) of \mathfrak{A} , and a unit vector $\Omega \in \mathcal{H}$ such that:*

1. $\omega(A) = \langle \Omega, \pi(A)\Omega \rangle$, for all $A \in \mathfrak{A}$;
2. $\pi(\mathfrak{A})\Omega$ is dense in \mathcal{H} .

Furthermore, the representation (\mathcal{H}, π) is the unique one (up to unitarily equivalence) satisfying the two conditions.

Since we will later need to invoke the details of the GNS construction, we sketch the outlines of its proof here.

Sketch of proof. We construct the Hilbert space \mathcal{H} from equivalence classes of elements in \mathfrak{A} , and the representation π is given by the action of left multiplication. In particular, define a bounded sesquilinear form on \mathfrak{A} by setting

$$\langle A, B \rangle_\omega = \omega(A^*B), \quad A, B \in \mathfrak{A}.$$

Let \mathcal{H}_0 be the quotient of \mathfrak{A} induced by the norm $\|A\|_\omega = \langle A, A \rangle_\omega^{1/2}$. Let \mathcal{H} be the unique completion of the pre-Hilbert space \mathcal{H}_0 . Thus there is an inclusion mapping $j : \mathfrak{A} \rightarrow \mathcal{H}$ with $j(\mathfrak{A})$ dense in \mathcal{H} . Define the operator $\pi(A)$ on \mathcal{H} by setting

$$\pi(A)j(B) = j(AB), \quad B \in \mathfrak{A}.$$

One must verify that $\pi(A)$ is well-defined, and extends uniquely to a bounded linear operator on \mathcal{H} . One must also then verify that π is a $*$ -homomorphism. Finally, if we let $\Omega = j(I)$, then Ω is obviously cyclic for $\pi(\mathfrak{A})$. ■

PROPOSITION 18. *Let ω be a state of \mathfrak{A} . The GNS representation (\mathcal{H}, π) of \mathfrak{A} induced by ω is irreducible iff ω is pure.*

Notes: Standard references on C^* -algebras include [Kadison and Ringrose, 1997] and [Takesaki, 2002].

1.3 Type classification of von Neumann algebras

DEFINITION 19. Two projections E, F in a von Neumann algebra \mathfrak{R} are said to be *equivalent*, written $E \sim F$ just in case there is a $V \in \mathfrak{R}$ such that $V^*V = E$ and $VV^* = F$.

NOTE 20. If we were being really careful, we would replace “equivalent” in the previous definition with “equivalence modulo \mathfrak{R} ”, and similarly “ \sim ” with “ $\sim_{\mathfrak{R}}$.” But we will not run into trouble by omitting the reference to \mathfrak{R} . The operator V in the previous definition is called a *partial isometry* with *initial projection* E and *final projection* F .

DEFINITION 21. For von Neumann algebras \mathfrak{R}_1 and \mathfrak{R}_2 , we let $\mathfrak{R}_1 \wedge \mathfrak{R}_2 = \mathfrak{R}_1 \cap \mathfrak{R}_2$. We let $\mathfrak{R}_1 \vee \mathfrak{R}_2$ denote the von Neumann algebra generated by \mathfrak{R}_1 and \mathfrak{R}_2 , i.e. the intersection of all von Neumann algebras containing \mathfrak{R}_1 and \mathfrak{R}_2 .

DEFINITION 22. $Z(\mathfrak{R}) = \mathfrak{R} \wedge \mathfrak{R}'$ is called the *center* of the von Neumann algebra \mathfrak{R} . A von Neumann algebra \mathfrak{R} is called a *factor* just in case $Z(\mathfrak{R}) = \mathbb{C}I$, equivalently, $\mathfrak{R} \vee \mathfrak{R}' = \mathfrak{B}(\mathcal{H})$. A projection $E \in Z(\mathfrak{R})$ is called a *central projection* in \mathfrak{R} .

DEFINITION 23. Let $E \in \mathfrak{A}$ be a projection, and let $E\mathfrak{A}E = \{EAE : A \in \mathfrak{A}\}$. Then clearly, $E\mathfrak{A}E$ is a linear subspace of \mathfrak{A} . Furthermore, since for $A, B \in \mathfrak{A}$, $AEB \in \mathfrak{A}$ and $A^* \in \mathfrak{A}$, it follows that $E\mathfrak{A}E$ is closed under products, as well as under $*$. It is also not difficult to see that $E\mathfrak{A}E$ is weakly closed, and hence is a von Neumann algebra on $E\mathcal{H}$.

DEFINITION 24. Let \mathfrak{A} be a von Neumann algebra. A projection $E \in \mathfrak{A}$ is said to be:

1. *minimal* just in case \mathfrak{A} contains no proper subprojection of E .
2. *abelian* just in case the algebra $E\mathfrak{A}E$ is abelian.
3. *infinite* just in case there is a projection $E_0 \in \mathfrak{A}$ such that $E_0 < E$ and $E \sim E_0$.
4. *finite* just in case it is not infinite.
5. *properly infinite* just in case E is infinite and for each central projection P of \mathfrak{A} , either $PE = 0$ or PE is infinite.

FACT 25. We have the following relations for projections:

$$\begin{aligned} \text{minimal} &\implies \text{abelian} \implies \text{finite} \\ \text{properly infinite} &\implies \text{infinite} \iff \neg \text{finite} \end{aligned}$$

For factors, the first arrows on both lines can be reversed.

We now give the Murray-von Neumann type classification of factors (for more on this, see [Kadison and Ringrose, 1997, Chap. 7] or [Sunder, 1987, Chap. 1]).

DEFINITION 26. A von Neumann factor \mathfrak{A} is said to be:

1. type I if it contains an abelian projection;
2. type II if it contains a finite projection, but no abelian projection;
3. type III if it is neither type I nor type II.

The type I factors were already completely classified by Murray and von Neumann: for each cardinal number κ there is a unique (up to isomorphism) type I_κ factor, namely $\mathfrak{B}(\mathcal{H})$ where \mathcal{H} is a Hilbert space of dimension κ . The type II factors can be further subdivided according to whether or not the identity projection I is finite (type II_1) or infinite (type II_∞). The type III factors can be subdivided into types III_λ with $\lambda \in [0, 1]$, although the basis for this subclassification depends on Tomita-Takesaki modular theory (see Section 1.4).

For general von Neumann algebras, the type classification must be just a bit more sophisticated: a type I algebra is defined as an algebra that has an abelian projection E such that no nontrivial projection in $Z(\mathfrak{A})$ majorizes E . Similarly, a

type II algebra is defined as an algebra having a finite projection E such that no nontrivial projection in $Z(\mathfrak{R})$ majorizes E . Thus we have:

PROPOSITION 27. *Let \mathfrak{R} be a von Neumann algebra. Then $\mathfrak{R} = \mathfrak{R}_I \oplus \mathfrak{R}_{II} \oplus \mathfrak{R}_{III}$, where \mathfrak{R}_X is type X for $X=I,II,III$.*

Proof. See [Kadison and Ringrose, 1997, Thm. 6.5.2]. ■

We will soon see that the local algebras in QFT are “typically” type III, and this has many interesting implications. The fact that type III algebras do not have abelian projections is connected to questions of locality in Section 3.3. The fact that the state space of type III₁ factors is homogeneous is also connected to questions of locality in Section 3.3. The fact that type III algebras do not contain representatives of their states (i.e. density operators) is connected to the modal interpretation of QFT in Section 5.

The following classification of von Neumann algebras is also natural, but it cuts across the Murray-von Neumann classification.

DEFINITION 28. A von Neumann algebra \mathfrak{R} is said to be:

- of *infinite type* if I is infinite in \mathfrak{R} ;
- *properly infinite* if I is properly infinite in \mathfrak{R} .
- *semi-finite* if the central projection E_{III} in \mathfrak{R} (defined in Prop. 27) is zero.

The finite factors include the type I_n and type II_1 factors. The infinite factors include the type I_∞ factors as well as the type II_∞ and III factors. The distinction between finite and infinite factors coincides with the existence of a tracial state.

DEFINITION 29. A faithful normalized trace on a von Neumann algebra \mathfrak{R} is a state ρ on \mathfrak{R} such that:

1. ρ is tracial; i.e. $\rho(AB) = \rho(BA)$, for all $A, B \in \mathfrak{R}$,
2. ρ is faithful; i.e. $\rho(A^*A) = 0$ only if $A = 0$.

FACT 30. A von Neumann factor \mathfrak{R} is finite iff there is a faithful normal tracial state ρ on \mathfrak{R} . A von Neumann factor \mathfrak{R} is semifinite iff there is a “faithful normal semifinite trace” on \mathfrak{R} ; but we do not pause here to define this notion.

1.4 Modular theory

We state here without proof some of the basic facts about Tomita-Takesaki modular theory. These facts are necessary in order to understand the classification of type III von Neumann algebras, which in turn is essential to understanding the mathematical structure of AQFT.

DEFINITION 31. Let \mathfrak{A} be a von Neumann algebra acting on a Hilbert space \mathcal{H} , and suppose that $\Omega \in \mathcal{H}$ is cyclic and separating for \mathfrak{A} . In such a case, we say that (\mathfrak{A}, Ω) is *in standard form*. Define an operator S_0 on \mathcal{H} by setting

$$S_0 A \Omega = A^* \Omega, \quad A \in \mathfrak{A}.$$

Then S_0 extends to a closed anti-linear operator S on \mathcal{H} . Let $S = J\Delta^{1/2}$ be the polar decomposition of S , so that Δ is positive (but generally unbounded), and J is anti-unitary. (Recall that a positive operator has spectrum in \mathbb{R}^+ .) We call Δ the *modular operator* and J the *modular conjugation* associated with the pair (\mathfrak{A}, Ω) .

THEOREM 32. *Let \mathfrak{A} be a von Neumann algebra with cyclic and separating vector Ω . Then $J\Omega = \Omega = \Delta\Omega$ and*

$$\begin{aligned} \Delta^{it}\mathfrak{A}\Delta^{-it} &= \mathfrak{A}, & \forall t \in \mathbb{R}, \\ J\mathfrak{A}J &= \mathfrak{A}'. \end{aligned}$$

Proof. See [Kadison and Ringrose, 1997, Thm. 9.2.9], or [Sunder, 1987, Thm. 2.3.3]. ■

DEFINITION 33. Let (\mathfrak{A}, Ω) be in standard form, and let ω be the state of \mathfrak{A} induced by Ω . For each $t \in \mathbb{R}$, define the *modular automorphism* σ_t^ω of \mathfrak{A} by

$$\sigma_t^\omega(A) = \Delta^{it} A \Delta^{-it}, \quad A \in \mathfrak{A},$$

for all $A \in \mathfrak{A}$. Define a $*$ anti-isomorphism $\gamma : \mathfrak{A} \rightarrow \mathfrak{A}'$ by setting $\gamma(A) = JA^*J$, for all $A \in \mathfrak{A}$.

DEFINITION 34. If \mathfrak{A} is a C^* -algebra, we let $\text{Inn}\mathfrak{A}$ denote the group of inner automorphisms of \mathfrak{A} ; i.e. $\alpha \in \text{Inn}\mathfrak{A}$ just in case there is a unitary $U \in \mathfrak{A}$ such that $\alpha(A) = UAU^*$ for all $A \in \mathfrak{A}$.

The spectrum of the modular operator Δ gives a rough measure of the periodicity of the modular automorphism group $(\sigma_t^\omega)_{t \in \mathbb{R}}$; i.e. the smaller the spectrum of Δ , the closer the automorphism σ_t^ω is to the identity $\iota : \mathfrak{A} \rightarrow \mathfrak{A}$. In the extreme case, if $\text{sp}\Delta = \{1\}$, then $\sigma_t^\omega = \iota$ for all $t \in \mathbb{R}$. Conversely, as Δ goes up to \mathbb{R}^+ , the group $(\sigma_t^\omega)_{t \in \mathbb{R}}$ tends toward being ergodic (i.e. having no fixed points).

DEFINITION 35. Define the *modular spectrum* $S(\mathfrak{A})$ of \mathfrak{A} by

$$S(\mathfrak{A}) = \bigcap_{\omega} \text{sp}(\Delta_{\omega}),$$

where ω runs over the family of faithful normal states of \mathfrak{A} , and Δ_{ω} are the corresponding modular operators.

PROPOSITION 36. *Let \mathfrak{A} be a von Neumann factor with cyclic and separating vector Ω . Then the following are equivalent:*

1. \mathfrak{A} is semifinite.
2. For all $t \in \mathbb{R}$, the modular automorphism σ_t^ω is inner; i.e. there is a unitary $U \in \mathfrak{A}$ such that $\sigma_t^\omega(A) = UAU^*$ for all $A \in \mathfrak{A}$.
3. $S(\mathfrak{A}) = \{1\}$.

Proof. See [Takesaki, 2003, p. 122] and [Sunder, 1987, p. 111]. ■

We now proceed to Connes’ subclassification of the type III factors. This subclassification uses the notion of the “period of the flow of weights” (where a weight is a generalization of the notion of a state). However, in order to bypass some background material, we use the following (provably equivalent) definition.

DEFINITION 37. A factor \mathfrak{A} of type III is said to be:

1. Type III₀ if $S(\mathfrak{A}) = \{0, 1\}$.
2. Type III _{λ} , $\lambda \in (0, 1)$, if $S(\mathfrak{A}) = \{\lambda^n : n \in \mathbb{Z}\} \cup \{0\}$.
3. Type III₁ if $S(\mathfrak{A}) = \mathbb{R}^+$.

The conditions in Defn. 37 do not bear their physical interpretation on their sleeve. That is, it is not immediately clear how the physics of type III _{λ} algebras differs (if at all) from that of type III _{μ} algebras, for $\lambda \neq \mu$. However, a result of Connes and Størmer [1978] cashes out some of the significance of the distinctions between different types of factors.

DEFINITION 38. Let \mathfrak{A} be a von Neumann algebra, and let \mathfrak{A}_* be its normal state space. We define the *diameter of the state orbit space* $d(\mathfrak{A})$ by

$$d(\mathfrak{A}) = \sup \left\{ \inf \{ \|(\omega_1 \circ \alpha) - \omega_2\| : \alpha \in \text{Inn}\mathfrak{A} \} : \omega_1, \omega_2 \in \mathfrak{A}_* \right\}.$$

Alternatively, let $[\omega]$ denote the norm closure of $\{\omega \circ \alpha : \alpha \in \text{Inn}\mathfrak{A}\}$ (the orbit of the state under inner automorphisms), and let K denote the quotient of the normal state space \mathfrak{A}_* . Then $d(\mathfrak{A})$ is the diameter of K relative to the induced metric

$$\bar{d}([\omega_1], [\omega_2]) = \inf \{ \|\omega'_1 - \omega'_2\| : \omega'_i \in [\omega_i] \}.$$

Clearly $d(\mathfrak{A}) \in [0, 2]$, with $d(\mathfrak{A}) = 0$ iff the orbit of every state is dense in the normal state space. If \mathfrak{A} is not a factor, then there are states ω_1, ω_2 such that $\|\omega_1 \circ \alpha - \omega_2\| = 2$ for all $\alpha \in \text{Inn}\mathfrak{A}$, and so $d(\mathfrak{A}) = 2$. For type I _{n} factors, the distance between normal states is the same as the trace norm distance of the corresponding density operators. In this case, we have

$$d(\mathfrak{A}) = 2 \left(1 - \frac{1}{n} \right) = \|\tau - \omega\|,$$

where τ is the trace and ω is any pure state. We also have $d(\mathfrak{K}) = 2$ for factors of type I_∞ and of type II [Takesaki, 2003, p. 430].

If $d(\mathfrak{K})$ gives some sort of measure of “how noncommutative” the algebra \mathfrak{K} is, then type III_1 factors are the most noncommutative.

DEFINITION 39. A von Neumann algebra \mathfrak{K} is said to be *countably decomposable* just in case any family of mutually orthogonal projection operators in \mathfrak{K} is countable.

PROPOSITION 40. *If \mathfrak{K} is a countably decomposable factor of type III_λ , then*

$$d(\mathfrak{K}) = 2 \frac{1 - \lambda^{1/2}}{1 + \lambda^{1/2}}.$$

Proof. See [Connes and Størmer, 1978] and [Takesaki, 2003, p. 427]. ■

The function $f(\lambda) = 2(1 - \lambda^{1/2})/(1 + \lambda^{1/2})$ is monotonically decreasing on $[0, 1]$. In particular, $f(1) = 0$ so that, for type III_1 factors, the orbit of any normal state ω is norm dense in the state space. According to Connes [1994, p. 473] this means that “one cannot distinguish between two states of a factor of type III_1 by means of a property that is closed and invariant under inner automorphisms.” In particular, since two unitarily equivalent states must be considered to be “equally mixed,” there are no distinctions to be drawn in terms of the mixedness of states of a type III_1 factor.

Notes: For an overview of modular theory, see [Summers, ND] or [Connes, 1994]. For a full treatment, see [Takesaki, 2003]. For a detailed exposition of applications of modular theory in QFT, see [Borchers, 2000].

2 STRUCTURE OF THE NET OF OBSERVABLE ALGEBRAS

2.1 Nets of algebras, basic properties

AQFT proceeds by isolating some structural assumptions that hold in most known QFT models. It formalizes these structural assumptions, and then uses “abstract but efficient nonsense” to derive consequences of these assumptions.

The basic formalism of AQFT is a “net of local observable algebras” over spacetime. Although this formalism can be applied to a very wide class of spacetimes, we restrict attention in this Chapter mostly to Minkowski spacetime.

An open *double cone* in Minkowski spacetime is the intersection of the causal future of a point x with the causal past of a point y to the future of x . Let \mathcal{K} be the set of open double cones in Minkowski spacetime, and let $O \mapsto \mathfrak{A}(O)$ be a mapping from \mathcal{K} to C^* -algebras. We assume that all our C^* -algebras are unital, i.e. have a multiplicative identity. We assume that the set $\{\mathfrak{A}(O) : O \in \mathcal{K}\}$ of C^* -algebras (called a *net of observable algebras over Minkowski spacetime*) is an inductive system in the sense that:

If $O_1 \subseteq O_2$, then there is an embedding (i.e. an isometric $*$ -homomorphism) $\alpha_{12} : \mathfrak{A}(O_1) \rightarrow \mathfrak{A}(O_2)$.

ASSUMPTION 41 (Isotony). The mapping $O \mapsto \mathfrak{A}(O)$ is an inductive system.

The isotony assumption is sometimes motivated by the idea that an observable measurable in a region O_1 is *a fortiori* measurable in any region O_2 containing O_1 . But the isotony axiom is also justified by its utility: for, if $\{\mathfrak{A}(O) : O \in \mathcal{K}\}$ is an inductive system, then there is an inductive limit C^* -algebra \mathfrak{A} generated by all the local algebras. We call \mathfrak{A} the *quasilocal* algebra, because it contains observables that can be uniformly approximated by local observables.

NOTE 42. In some spacetimes, the set of double cones is not directed. In many such cases, it is still possible to define the quasilocal algebra by means of more sophisticated techniques [Fredenhagen, 1993].

Now we turn to the main *relativistic* assumption of AQFT.

ASSUMPTION 43 (Microcausality). A net \mathfrak{A} of C^* -algebras is said to satisfy *microcausality* just in case if O_1, O_2 are spacelike separated double cones, then $[\mathfrak{A}(O_1), \mathfrak{A}(O_2)] = \{0\}$.

This assumption is thought to reflect the constraints on spacetime structure imposed by the theory of relativity.

NOTE 44. It is *not* a tenet of AQFT that quantities that are associated to spacelike separated regions must be represented by commuting operators. In fact, Fermi field operators assigned to spacelike separated regions will anticommute. So, AQFT has need of a distinction between observable (represented by elements of $\mathfrak{A}(O)$) and unobservable quantities (represented by “field operators”). For more on this distinction, see Sections 7.2 and following on DHR superselection theory.

In this Chapter, we will not attempt to justify or to dispute the microcausality assumption. However, we will briefly discuss its connection to issues of locality in Section 3.

2.2 Existence/uniqueness of vacuum states/representations

The existence of translation-invariant states

In this section, we inquire concerning the existence and uniqueness of vacuum states and representation. For this, recall that an *affine space* (e.g. Minkowski spacetime) is a triple consisting of a set S , a vector space V , and a map $+ : S \times V \rightarrow S$ satisfying certain properties. In this case, V is called the *translation group*.

ASSUMPTION 45 (Translation Covariance). If \mathfrak{A} is a net of operator algebras on an affine space, then we assume that there is a faithful, continuous representation $x \mapsto \alpha_x$ of the translation group in the group $\text{Aut}\mathfrak{A}$ of automorphisms of \mathfrak{A} , and

$$\alpha_x(\mathfrak{A}(O)) = \mathfrak{A}(O + x),$$

for any double cone O , and translation x .

NOTE 46. For the case of Minkowski spacetime, the translation group is a subgroup of the Poincaré group. In many cases of physical interest, $x \rightarrow \alpha_x$ extends to a representation of the full Poincaré group in the group $\text{Aut}\mathfrak{A}$ of automorphisms of \mathfrak{A} . But we will only need that fact for one result (Prop. 132).

Translation invariance has traditionally been thought to be a necessary condition on a vacuum state.

FACT 47. If there is an action α of the translation group on \mathfrak{A} , then translation-invariant states of \mathfrak{A} exist. Indeed, since the translation group is abelian, it has an invariant mean μ — i.e. a translation invariant, positive linear functional on the algebra $L^\infty(G)$ of essentially bounded measurable (with respect to the Haar measure) functions on the group G . Given a state ω of \mathfrak{A} , we can then define an averaged state ρ by

$$\rho(A) := \int \omega(\alpha_x A) d\mu(x).$$

The state ρ is translation invariant. (See Emch, this volume, Section 3.5.)

NOTE 48. The preceding argument cannot be used to show the existence of Lorentz invariant states. The Lorentz group is not *amenable*, and so does not admit an invariant mean. Hence, we cannot use these general methods to prove the existence of Lorentz invariant states. Of course, in concrete models (e.g. free Bose and Fermi fields) there are other way to establish the existence of such states.

Let G be a group acting by automorphisms on \mathfrak{A} . A generalization of the GNS theorem shows that a G -invariant state ω of \mathfrak{A} gives rise to a GNS Hilbert space \mathcal{H} that carries a unitary representation U of G , and the GNS vector Ω is invariant under the G -action on \mathcal{H} .

FACT 49. Let α be a strongly continuous action of the group G by automorphisms of \mathfrak{A} . If ω is a G -invariant state of \mathfrak{A} , then the GNS representation (\mathcal{H}, π) of \mathfrak{A} induced by ω is G -covariant in the sense that there is a strongly continuous representation U of G in the unitary group of $\mathfrak{B}(\mathcal{H})$ such that

1. $U(g)\pi(A)U(g)^* = \pi(\alpha_g(A))$, for all $A \in \mathfrak{A}$,
2. $U(g)\Omega = \Omega$ for all $g \in G$.

Only one vacuum per Hilbert space

NOTE 50. When considering the group $\text{Aut}\mathfrak{A}$ of automorphisms of a C^* -algebra, we take as our standard topology the strong topology on the set $L(\mathfrak{A})$ of bounded linear mappings on \mathfrak{A} (considered as a Banach space). That is, α_i converges to α just in case for each $A \in \mathfrak{A}$, $\alpha_i(A)$ converges to $\alpha(A)$ in the norm on \mathfrak{A} .

DEFINITION 51. We use the GNS representation theorem (Thm. 17) to transfer terminology about representations (Defn. 14) to terminology about states. So, e.g., we say that two states are *disjoint* if their GNS representations are disjoint.

A vacuum state should be at least translation invariant. Furthermore, the microcausality assumption on the net \mathfrak{A} entails that any two observables commute “in the limit” where one is translated out to spacelike infinity. That is, for any $A, B \in \mathfrak{A}$, and for any spacelike vector x ,

$$\lim_{t \rightarrow \infty} \|[\alpha_{tx}(A), B]\| = 0.$$

This in turn entails that G acts on \mathfrak{A} as a *large group of automorphisms* in the following sense:

If ω is a G -invariant state and (\mathcal{H}, π) is the GNS representation of \mathfrak{A} induced by ω , then for any $A \in \mathfrak{A}$,

$$\overline{\text{conv}}\{\pi(\alpha_g(A)) : g \in G\},$$

has nonempty intersection with $\pi(\mathfrak{A})'$.

Here we use $\overline{\text{conv}}S$ to denote the weakly closed convex hull of S . (See [Størmer, 1970] for the relevant proofs.) Note however that we would also expect the same to be true in a non-relativistic setting, because we would expect observables associated with disjoint regions of space to commute. (We have not invoked the fact that any vector in Minkowski spacetime is the sum of two spacelike vectors.)

Thanks to extensive research on “ C^* -dynamical systems,” much is known about G -invariant states when G acts as a large group of automorphisms of \mathfrak{A} . In particular, the set of G -invariant states is convex and closed (in the weak* topology), hence the set has extreme points, called *extremal invariant states*. (Obviously if a pure state of \mathfrak{A} is G -invariant, then it is extremal invariant.) Furthermore, we also have the following result concerning the disjointness of G -invariant states.

PROPOSITION 52. *Let ω be a G -invariant state of \mathfrak{A} , let \mathcal{H} be its GNS Hilbert space, and let Ω be the GNS vector. Then the following are equivalent:*

1. ω is clustering in the sense that

$$\lim_{t \rightarrow \infty} \omega(\alpha_{tx}(A)B) = \omega(A)\omega(B).$$

2. ω is extremal invariant.

3. If a G -invariant state ρ is quasiequivalent to ω , then $\rho = \omega$. In other words, no other G -invariant state is quasiequivalent to ω .

4. The ray spanned by Ω is the unique (up to scalar multiples) G -invariant subspace of \mathcal{H} .

Proof. See [Størmer, 1970]. For related details, see also [Emch, 1972, pp. 183, 287] and Emch, this volume, Section 3. ■

So, if a (vacuum) state is clustering, then no other translation invariant state is in its folium (i.e. the set of states that are quasiequivalent to that state). Similarly, if a state is *extremal* invariant (*a fortiori* if it is pure) then it is the unique translation invariant state in its folium.

NOTE 53. The existence of disjoint vacua is related to spontaneous symmetry breaking. See Section 10.7.

NOTE 54. Prop. 52 plays a central role in the proof of “Haag’s theorem” given in [Emch, 1972, p. 248]. In particular, the uniqueness of extremal G -invariant states is equated with the nonexistence of “vacuum polarization.”

2.3 The Reeh-Schlieder Theorem

We have assumed that a vacuum state is translation invariant. But we expect a vacuum state to obey a stronger constraint that reflects the relativistic nature of the theory. In particular, the unitary representation defined in Fact 49 is generated infinitesimally by the four momentum operator \mathbf{P} . (The idea of a four momentum operator can be made precise in the “SNAG [Stone-Naimark-Ambrose-Gelfand] Theorem,” which generalizes Stone’s theorem on the existence of self-adjoint operators generating one-parameter unitary groups.) We require that the energy is positive in every Lorentz frame, equivalently, that the spectrum of \mathbf{P} lies in the forward light cone.

We now generalize this requirement by abstracting away from the details of the forward lightcone. The forward lightcone G_+ has the following property: $G_+ \cap (-G_+) = \{0\}$ where $-G_+ = \{-g : g \in G_+\}$. So, the spectrum condition only requires that the unitary representation of the translation group has spectrum in a set that is asymmetric under taking additive inverses.

ASSUMPTION 55 (Spectrum Condition). Let G be the translation group, and let ω be a G -invariant state of \mathfrak{A} . We say that the pair (\mathfrak{A}, ω) satisfies the *spectrum condition* just in case: there is a subset G_+ of G such that $G_+ \cap (-G_+) = \{0\}$, and in the GNS representation (\mathcal{H}, π) of \mathfrak{A} induced by ω , the spectrum $\text{sp}(U)$ of the induced unitary representation of G , is contained in G_+ .

The Reeh-Schlieder Theorem shows that the spectrum condition entails that the vacuum vector Ω is cyclic for *every* local algebra. For this theorem, we suppose that a translation invariant vacuum state ω on \mathfrak{A} has been chosen, and that (\mathcal{H}, π) is the GNS representation of \mathfrak{A} induced by ω . We then define a corresponding net \mathfrak{R} of von Neumann algebras on \mathcal{H} by

$$O \mapsto \mathfrak{R}(O) \equiv \pi(\mathfrak{A}(O))''.$$

If the net \mathfrak{A} satisfies microcausality, then so will \mathfrak{R} . Since Ω is cyclic for $\pi(\mathfrak{A})$, the set $\{\mathfrak{R}(O)\Omega : O \in \mathcal{K}\}$ is dense in \mathcal{H} .

To prove the theorem, we need one additional assumption.

ASSUMPTION 56. The net $O \mapsto \mathfrak{R}(O)$ is said to satisfy *additivity* just in case for any double cone O , the set $\{\mathfrak{R}(O+x) : x \in G\}$ generates \mathfrak{R} as a C^* -algebra. (Here again, G denotes the translation group.)

The additivity assumption is sometimes justified on the grounds that there should be no smallest length scale in the theory — i.e. any observable is generated by taking products, sums, etc. of observables from arbitrarily small regions.

THEOREM 57 (Reeh-Schlieder). *Suppose that the net $O \mapsto \mathfrak{R}(O)$ satisfies the spectrum condition and additivity. Then for all double cones O , Ω is cyclic for $\mathfrak{R}(O)$. If the net \mathfrak{R} also satisfies microcausality, then Ω is separating for every local algebra.*

The Reeh-Schlieder (RS) Theorem has been one of the more intensely studied issues in the foundations of relativistic QFT. In a pair of articles [Redhead, 1995a; Redhead, 1995b], Redhead shows that the RS Theorem entails that the vacuum state displays nonlocal correlations. (See also [Halvorson and Clifton, 2000]). Redhead also points out since the vacuum is separating for each local algebra, every local event has a nonzero probability of occurring in the vacuum state; in particular, there can be no local number operators (since they would have the vacuum state as an eigenvector). Finally, [Fleming, 2000] argues that RS Theorem entails a pernicious sort of nonlocality, worse than the nonlocality in non-relativistic QM, and so indicates a need to revise the standard formulation of AQFT. (For one possible reply, see [Halvorson, 2001].)

Due to the use of the spectrum condition, it would seem that RS Theorem is a “purely relativistic result,” without analogue in non-relativistic QM or QFT (see [Saunders, 1992]). Furthermore, we might expect that many other results of relativistic QFT that are derived from RS Theorem would fail for non-relativistic theories. Indeed, non-relativistic QFT *does* admit local number operators. However, a version of the spectrum condition, and consequently a version of RS Theorem has been shown to hold for non-relativistic theories [Requardt, 1986].

Notes: The original Reeh-Schlieder Theorem was formulated in the axiomatic approach to QFT, and can be found in [Reeh and Schlieder, 1961]. More up-to-date presentations of the theorem can be found in [Horuzhy, 1990; D’Antoni, 1990; Baumgärtel and Wollenberg, 1992], and [Araki, 1999].

2.4 The funnel property

DEFINITION 58. Let $\mathfrak{R}_1, \mathfrak{R}_2$ be von Neumann algebras on \mathcal{H} such that $\mathfrak{R}_1 \subseteq \mathfrak{R}_2$. If there is a vector $\Omega \in \mathcal{H}$ that is cyclic and separating for $\mathfrak{R}_1, \mathfrak{R}_2$, and $\mathfrak{R}'_1 \cap \mathfrak{R}_2$, then the pair $(\mathfrak{R}_1, \mathfrak{R}_2)$ is said to be a *standard inclusion* of von Neumann algebras.

NOTE 59. Let $O \mapsto \mathfrak{R}(O)$ be a net of von Neumann algebras on Minkowski spacetime. Suppose that the Reeh-Schlieder property holds for Ω , i.e. for each

double cone O , Ω is cyclic and separating for $\mathfrak{A}(O)$. Then if O_1, O_2 are double cones such that the closure $\overline{O_1}$ of O_1 is contained in O_2 , then the pair $(\mathfrak{A}(O_1), \mathfrak{A}(O_2))$ is a standard inclusion of von Neumann algebras.

DEFINITION 60. Let $\mathfrak{A}_1, \mathfrak{A}_2$ be von Neumann algebras on \mathcal{H} such that $\mathfrak{A}_1 \subseteq \mathfrak{A}_2$. The pair $(\mathfrak{A}_1, \mathfrak{A}_2)$ is said to be a *split inclusion* if there is a type I factor \mathfrak{N} such that $\mathfrak{A}_1 \subseteq \mathfrak{N} \subseteq \mathfrak{A}_2$.

ASSUMPTION 61 (Funnel Property). The net $O \mapsto \mathfrak{A}(O)$ of von Neumann algebras is said to satisfy the *funnel property* if for any double cones O_1, O_2 with $\overline{O_1}$ contained in O_2 , the pair $(\mathfrak{A}(O_1), \mathfrak{A}(O_2))$ is a split inclusion.

NOTE 62. A type I factor \mathfrak{N} is countably decomposable iff \mathfrak{N} is isomorphic to $\mathfrak{B}(\mathcal{H})$ with \mathcal{H} separable iff \mathfrak{N} is separable in the ultraweak topology (see [Kadison and Ringrose, 1997, Exercise 5.7.7]).

In our discussion of superselection theory (Sections 7.2–11), at one crucial juncture (Prop. 243, p. 813) we will have to invoke the assumption that the vacuum Hilbert space is separable. This will be the only place in the Chapter where we need to assume that a Hilbert space is separable. In particular, the separability assumption is needed to establish the correspondence between two notions of superselection sectors, one of which is physically motivated, and one of which is mathematically useful. The following result is the only attempt we will make to connect the separability assumption to something with (perhaps) more clear physical significance. (In general, we are highly suspicious of the physical warrant for the separability assumption; compare with Section 6, and with [Halvorson, 2004].)

PROPOSITION 63. *Let \mathfrak{A} be a net of von Neumann algebras on \mathcal{H} , and suppose that $\Omega \in \mathcal{H}$ is cyclic and separating for all local algebras. If the net satisfies the funnel property, then \mathcal{H} is separable.*

Proof. (Compare with Prop. 1.6 of [Doplicher and Longo, 1984].) Let O_1, O_2 be double cones with $\overline{O_1} \subseteq O_2$. Let \mathfrak{N} be a type I factor such that $\mathfrak{A}(O_1) \subseteq \mathfrak{N} \subseteq \mathfrak{A}(O_2)$, and let ω be the state of \mathfrak{N} induced by Ω . Recall that \mathfrak{N} is isomorphic to $\mathfrak{B}(\mathcal{K})$ for some Hilbert space \mathcal{K} . Since $\mathfrak{N} \subseteq \mathfrak{A}(O_2)$ and Ω is separating for $\mathfrak{A}(O_2)$, ω is faithful and normal. Hence \mathcal{K} is separable, and there is a countable set \mathfrak{N}_0 that is ultraweakly dense in \mathfrak{N} . Since $\mathfrak{A}(O_1) \subseteq \mathfrak{N}$, and Ω is cyclic for $\mathfrak{A}(O_1)$ it follows that $[\mathfrak{N}_0\Omega] = [\mathfrak{N}\Omega] = \mathcal{H}$. Hence \mathcal{H} is separable. ■

If one wanted to justify an assumption that the vacuum Hilbert space is separable, Prop. 63 shows that it is enough to justify the funnel property. There are concrete models where the funnel property demonstrably does *not* hold [Horuzhy, 1990, p. 23]. But the physical significance of these models is not clear, and there are a couple of other considerations that might favor the funnel property: (i): In Section 3.3, we show that connection of the funnel property with issues about nonlocality. (ii): Buchholz and Wichmann [1986] argue that the funnel property is a *sufficient* condition for a particle interpretation of QFT. Of course, the interpreter of QFT will want to critically examine Buchholz and Wichmann’s notion

of a “particle interpretation.” (Compare with Section 4.5, where particle interpretations are discussed further. Compare also with Section 6.2, which hints at connections between nonseparable Hilbert space and field interpretations of QFT.)

NOTE 64. The funnel property for free fields is shown in [Buchholz, 1974].

2.5 Type of local algebras

We now collect the currently known information on the type of local algebras in physically relevant representations of the net of local observable algebras.

DEFINITION 65. Let \mathfrak{R}_1 and \mathfrak{R}_2 be nets of von Neumann algebras on a Hilbert space \mathcal{H} . We say that \mathfrak{R}_1 and \mathfrak{R}_2 are *locally quasiequivalent* just in case for each double cone O there is an isomorphism $\varphi_O : \mathfrak{R}_1(O) \rightarrow \mathfrak{R}_2(O)$.

NOTE 66. Although it is not an “axiom” of AQFT, there are good reasons to believe that representations of physical interest (in particular for elementary particle physics) are locally quasiequivalent to some vacuum representation, where a vacuum representation is the GNS representation of some privileged (e.g. perhaps translation invariant) state. For example local quasi-equivalence holds between any two physical representations according to the selection criterion of Doplicher-Haag-Roberts (see Section 7.2 and following), and according to the more liberal selection criterion of [Buchholz and Fredenhagen, 1982]. Thus, any conclusion we draw concerning the structure of *local algebras* in a vacuum representation can be inferred to hold as well for these other representations.

Local algebras are properly infinite

Some relatively simple results narrow down the possible options for the type of local algebras. For this, we define the important “property B,” because it is a consequence of plausible assumptions (viz. additivity and the spectrum condition), because it also makes sense in situations where there is no translation group (unlike the spectrum condition), and because it is all we need to infer various results, in particular that local algebras are properly infinite.

DEFINITION 67. Let $O \rightarrow \mathfrak{R}(O)$ be a net of von Neumann algebras on some Hilbert space \mathcal{H} . We say that the net \mathfrak{R} satisfies *property B* just in case for any two double cones O_1 and O_2 such that $\overline{O_1} \subseteq O_2$, if $E \in \mathfrak{R}(O_1)$ is a nonzero projection, then E is equivalent in $\mathfrak{R}(O_2)$ to the identity projection I ; i.e. there is an isometry $V \in \mathfrak{R}(O_2)$ such that $VV^* = E$.

NOTE 68. If for each O , the algebra $\mathfrak{R}(O)$ is type III, then the net \mathfrak{R} satisfies property B.

We expect property B to hold for a net of observable algebras because it follows from the physically motivated postulates of weak additivity and the spectrum condition.

PROPOSITION 69. *Let $O \mapsto \mathfrak{K}(O)$ be a net of von Neumann algebras satisfying microcausality, the spectrum condition, and weak additivity. Then the net $O \mapsto \mathfrak{K}(O)$ satisfies property B.*

Proof. For the original proof, see [Borchers, 1967]. For a recent exposition, see [D’Antoni, 1990]. ■

ASSUMPTION 70 (Nontriviality). A net $O \mapsto \mathfrak{A}(O)$ of C^* -algebras is said to satisfy *non-triviality* just in case for each double cone O , $\mathfrak{A}(O) \neq \mathbb{C}I$.

PROPOSITION 71. *Let $O \mapsto \mathfrak{K}(O)$ be a net of von Neumann algebras that satisfies microcausality, property B, and non-triviality. Then for every double cone O , the von Neumann algebras $\mathfrak{K}(O)$ and $\mathfrak{K}(O)'$ are properly infinite.*

Proof. We first show that $\mathfrak{K}(O)$ is properly infinite; that is, that every central projection in $\mathfrak{K}(O)$ is infinite. Let C be a central projection in $\mathfrak{K}(O)$. Choose a nontrivial double cone O_1 whose closure is contained in O . Then by property B, for each nonzero projection $E \in \mathfrak{K}(O_1)$, E is equivalent to I modulo $\mathfrak{K}(O)$. Since $\mathfrak{K}(O_1) \neq \mathbb{C}I$, there is a projection $E \in \mathfrak{K}(O_1)$ such that $E \sim (I - E) \sim I$ modulo $\mathfrak{K}(O)$. It then follows that $EC \sim (I - E)C \sim C$ modulo $\mathfrak{K}(O)$. It is clear that $EC = CEC \leq C$. If $EC = C$ then $(I - E)C = 0$, a contradiction. Therefore $EC < C$ and $EC \sim C$ modulo $\mathfrak{K}(O)$. That is, C is an infinite projection in $\mathfrak{K}(O)$, and $\mathfrak{K}(O)$ is properly infinite. By microcausality, $\mathfrak{K}(O_1) \subseteq \mathfrak{K}(O)'$; thus the preceding argument also shows that $\mathfrak{K}(O)'$ is properly infinite. ■

In particular, the preceding proposition rules out the cases of type I_n and type II_1 von Neumann algebras. Already this result has implications for questions about nonlocality; see Prop. 102 in Section 3.3. However, the previous proposition leaves open the possibility that local algebras might be type I_∞ factors, and it also leaves open the case that local algebras might be direct sums of heterogeneous types of von Neumann algebras.

Local algebras are hyperfinite

We will shortly see that the best results we have point toward the fact that local algebras are type III, which were originally thought to be unruly anomalies with no relevance for physics. However, we first show that under some physically plausible conditions, local algebras are approximated by finite-dimensional algebras (i.e. they are “hyperfinite”), which shows that after all they are not so unruly.

DEFINITION 72. Let \mathfrak{K} be a von Neumann algebra. Then \mathfrak{K} is said to be *hyperfinite* just in case there is a family $(\mathfrak{K}_a)_{a \in \mathbb{A}}$ of finite dimensional von Neumann algebras in \mathfrak{K} such that $\mathfrak{K} = (\cup_{a \in \mathbb{A}} \mathfrak{K}_a)''$.

Hyperfiniteness turns out to be an extremely useful condition for mathematical purposes. Indeed, hyperfiniteness is intimately linked to the existence of normal conditional expectations (see [Kadison and Ringrose, 1997, Chap. 8]), and there

is a unique type II₁ hyperfinite factor, and a unique type III₁ hyperfinite factor. From a physical/foundational point of view, one might also think that a failure of hyperfiniteness for \mathfrak{R} might make it difficult to find a correspondence between elements of the algebra \mathfrak{R} and real-life laboratory procedures which can only involve a finite number of tasks.

FACT 73. Every type I von Neumann algebra is hyperfinite. See [Kadison and Ringrose, 1997, Exercise 8.7.26].

ASSUMPTION 74 (Inner/Outer Continuity). A net $O \mapsto \mathfrak{R}(O)$ of von Neumann algebras is said to be *inner continuous* if for any monotonically increasing net $(O_a)_{a \in \mathbb{A}}$ with least upper bound O , we have

$$\bigvee_{a \in \mathbb{A}} \mathfrak{R}(O_a) = \mathfrak{R}(O),$$

where $\mathfrak{R}_1 \vee \mathfrak{R}_2$ denotes the von Neumann algebra generated by \mathfrak{R}_1 and \mathfrak{R}_2 . Outer continuity is defined by taking a decreasing net of regions, and the intersection of the corresponding von Neumann algebras.

NOTE 75. The condition that the net \mathfrak{R} be continuous from the inside is satisfied whenever \mathfrak{R} is the “minimal” net constructed in the standard way from underlying Wightman fields. See [Buchholz *et al.*, 1987]. Similarly, the maximal net satisfies outer continuity.

PROPOSITION 76. *Suppose that the net $O \mapsto \mathfrak{R}(O)$ satisfies the funnel property and either inner or outer continuity. Then for each double cone O , $\mathfrak{R}(O)$ is hyperfinite.*

Sketch of proof. (Compare [Buchholz *et al.*, 1987, p. 134].) We just look at the case where the net is inner continuous. By the funnel property there is a type I factor \mathfrak{N}_i interpolating between $\mathfrak{R}(O_i)$ and $\mathfrak{R}(O)$. It then follows that the union of the ascending sequence \mathfrak{N}_i of hyperfinite factors is dense in $\mathfrak{R}(O)$, hence $\mathfrak{R}(O)$ is hyperfinite. ■

Local algebras are type III₁ factors

A series of results, accumulated over a period of more than thirty years, indicates that the local algebras of relativistic QFT are type III von Neumann algebras, and more specifically, hyperfinite type III₁ factors. We gather some of these results in this section. The first result, due to Longo [1979], improved on some earlier results by Driessler.

PROPOSITION 77. *Let \mathfrak{R} be a von Neumann algebra acting on \mathcal{H} , $\Omega \in \mathcal{H}$ a separating unit vector for \mathfrak{R} , G a locally compact abelian group with dual Γ , and U a continuous unitary representation of G on \mathcal{H} such that $U\Omega = \Omega$ and the ray $\mathbb{C}\Omega$ is the unique $U(G)$ invariant subspace of \mathcal{H} . Suppose that there exist subsets $G_+ \subseteq G$ and $\Gamma_+ \subseteq \Gamma$ such that*

1. $G_+ \cup (-G_+) = G$ and $U(g)\mathfrak{R}U(g)^* \subseteq \mathfrak{R}$, for all $g \in G_+$.
2. $\Gamma_+ \cap (-\Gamma_+) = \{0\}$ and $\text{sp}(U) \subseteq \Gamma_+$.

Then either $\mathfrak{R} = \mathbb{C}I$ or \mathfrak{R} is a type III₁ factor.

Sketch of proof. (See [Longo, 1979, p. 203] for details.) Let ω be the state of \mathfrak{R} given by $\omega(A) = \langle \Omega, A\Omega \rangle$. The proof of this result proceeds by showing that $\mathfrak{R}_\omega = \mathbb{C}I$, where \mathfrak{R}_ω is the centralizer of the state ω . In particular, let E be a projection in \mathfrak{R}_ω , and define the function $f : G \rightarrow \mathbb{C}$ by

$$f(g) = \langle \Omega, EU(g)E\Omega \rangle = \langle \Omega, EU(g)EU(-g)\Omega \rangle.$$

Using the constraint on $\text{sp}(U)$, it can be shown that f is constant, and hence $U(g)E\Omega = E\Omega$ for all $g \in G$. Since $\mathbb{C}\Omega$ is the unique invariant subspace under $U(G)$, it follows that $E\Omega = \Omega$, and since Ω is separating for \mathfrak{R} , $E = 0$ or $E = I$. ■

The preceding proposition applies to algebras of the form $\pi(\mathfrak{A}(W))''$, where W is a wedge region, and π is a vacuum representation of the quasilocal algebra \mathfrak{A} . Indeed, we can take G_+ to be a one-parameter semi-group of lightlike translations with origin at the apex of W , in which case $\mathbb{R} = G_+ \cup (-G_+)$. Let ω be a translation invariant state on \mathfrak{A} such that (\mathfrak{A}, ω) satisfies that spectrum condition (Assumption 55). We then have that the dual group Γ of G in \mathbb{R}^4 is also a lightlike line, and hence the spectrum condition entails that there is a subset Γ_+ of Γ , namely those vectors that point toward the future, such that $\Gamma_+ \cap (-\Gamma_+) = \{0\}$. Finally, we saw in Section 2.2 that when ω is extremal invariant, the ray $\mathbb{C}\Omega$ is the unique $U(G)$ invariant subspace of \mathcal{H} .

For results relevant to *local* algebras, we must impose one further condition on the net \mathfrak{R} . The first result ([Buchholz *et al.*, 1987]) requires reference to axiomatic QFT with unbounded operators smeared by test-functions (see [Streater and Wightman, 1964]). That is, we must assume that the net \mathfrak{R} arises from an underlying Wightman field theory that satisfies a certain condition — asymptotic scale invariance.

Recall that in the axiomatic approach, fields are essentially self-adjoint operators of the form $\Phi(f)$, where f is a test-function on spacetime. The presence of these test-functions allows the definition of a notion of asymptotic scale invariance.

DEFINITION 78. Let $N : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be a monotone function. Then a scaling transformation of the test-functions is given by $f \mapsto f_\lambda$, where $f_\lambda(x) = N(\lambda)f(\lambda^{-1}x)$. Let Φ_α be a set of Wightman fields generating the net $O \mapsto \mathfrak{R}(O)$. We say that the fields satisfy *asymptotic scale invariance* just in case there is some field Φ with vanishing vacuum expectation values:

$$\langle \Omega, \Phi(f)\Omega \rangle = 0,$$

and for a suitable choice of $N(\lambda)$, the scaled field operators $\Phi(f_\lambda)$ have the following properties:

1. The expectation values $\langle \Omega, \Phi(f_\lambda)^* \Phi(f_\lambda) \Omega \rangle$ converge for all test-functions in the limit $\lambda \rightarrow 0$, and are nonzero for some f ;
2. The norms $\|\Phi(f_\lambda)^* \Phi(f_\lambda) \Omega\|$ and $\|\Phi(f_\lambda) \Phi(f_\lambda)^* \Omega\|$ stay bounded in this limit.

When a net of von Neumann algebras arises from a Wightman theory with asymptotic scale invariance, it follows that local algebras are hyperfinite type III₁ factors.

PROPOSITION 79 ([Buchholz *et al.*, 1987]). *Let \mathfrak{A} be a net of von Neumann algebras that satisfies microcausality, the spectrum condition, and the funnel property. Suppose also that \mathfrak{A} can be constructed from an underlying Wightman theory that satisfies asymptotic scale invariance. Then for each double cone O , $\mathfrak{A}(O) = \mathfrak{M} \bar{\otimes} \mathfrak{Z}$, where \mathfrak{M} is the unique type III₁ hyperfinite factor and \mathfrak{Z} is the center of $\mathfrak{A}(O)$.*

NOTE 80. In [Buchholz *et al.*, 1987], the funnel property is derived from a more basic postulate called “nuclearity,” which imposes bounds on the number of local degrees of freedom.

Of course, one wishes for a result that is more intrinsic to AQFT. Such a result is provided in [Buchholz and Verch, 1995], using the method of scaling algebras that allows the computation of the short distance (scaling) limit of a net \mathfrak{A} of local observables. (For a short exposition of scaling algebras, we refer the reader to [Buchholz, 1998].) In summary, besides the basic assumptions on the net, the only additional assumption needed to derive the type III₁ property is that the net has a nontrivial scaling limit.

NOTE 81. In some concrete models, it can be shown directly that local algebras are the unique type III₁ hyperfinite factor. For example, for the free Bose field of mass $m = 0$ (in the Minkowski vacuum representation), local algebras are isomorphic to algebras for wedge regions. Thus Prop. 77 shows that local algebras are type III₁ factors. Furthermore, the free Bose field of mass $m > 0$ is locally quasiequivalent to the case of $m = 0$, and so its local algebras are also type III₁ hyperfinite factors. See [Horuzhy, 1990, p. 254].

The derivation of the type III₁ property is one of the most surprising and interesting results of contemporary mathematical physics. But what is the foundational significance of the result? How would the world be different if local algebras were, say, type III_{1/2}, or even more radically different, if they were type II_∞? For one, there is a crucial difference between the structure of states on familiar type I algebras, and the structure of states on type III algebras: since type III algebras have no atomic projections, and the support projection of a pure normal state is atomic, it follows that type III algebras have no pure normal states. (But of course the same is true for type II algebras.) As pointed out in [Clifton and Halvorson, 2001b] and [Ruetsche, 2004], this absence of pure states is a further obstacle to an ignorance interpretation of quantum probabilities. (See also Section 3.4.)

Yngvason [2005] makes several interesting claims about the conceptual importance of type III algebras, especially in relation to questions of nonlocality. First, according to Yngvason, “type I intuitions” can lead to paradoxes, such as that encountered in Fermi’s famous two-atom system. However, claims Yngvason, these paradoxes disappear if we model these situation appropriately with type III algebras. Second, Yngvason claims that the homogeneity of the state space of a type III₁ factor \mathfrak{R} can be interpreted as saying that for any two states ω_1, ω_2 on \mathfrak{R} , ω_2 can be prepared from ω_1 (within arbitrarily good accuracy) via a unitary operation. Such an operation is, of course, nonselective, and so does not change the statistics of measurements of observables in \mathfrak{R}' . So, in one sense, an observer with a type III algebra has more control over his state space than an observer with a type I algebra.

3 NONLOCALITY AND OPEN SYSTEMS IN AQFT

NOTE 82. For this section, we use the following notational conventions: uppercase roman letters for algebras, lowercase roman letters for operators, and $\mathbb{1}$ for the multiplicative identity in an algebra.

It is a basic assumption of AQFT that the observable algebras $A(O_1)$ and $A(O_2)$ are mutually commuting when O_1 and O_2 are spacelike separated. This requirement — which we have called “microcausality” — is sometimes also called “Einstein causality,” because of a suggested connection between the commutativity of the algebras $A(O_1), A(O_2)$ and the relativistic prohibition on “superluminal signaling.” Implicit in this connection is a claim that if $[a, b] \neq 0$ for $a \in A(O_1)$ and $b \in A(O_2)$, then a measurement of a could change the statistics of a measurement of b .

Despite the fact that nonrelativistic QM makes no reference to spacetime, it has a footprint of the relativistic prohibition of superluminal signalling. In particular, the state space of two distinct objects is a tensor product $H_1 \otimes H_2$, and their joint algebra of observables is $B(H_1) \otimes B(H_2)$. In this tensor product construction we represent observables for system A as simple tensors $a \otimes \mathbb{1}$ and observables of system B as $\mathbb{1} \otimes b$. Thus, we have a version of microcausality. But we also have stronger independence properties. For example, for every state φ_1 of system A and state φ_2 of system B , there is a state φ of $A \otimes B$ such that $\varphi|_A = \varphi_1$ and $\varphi|_B = \varphi_2$.

In this section, we investigate the extent to which two local algebras $A(O_1), A(O_2)$ can be thought to represent distinct, independent parts of reality. In Sections 3.1 and 3.2, we discuss the relations between microcausality and other independence assumptions for the algebras $A(O_1), A(O_2)$. In Section 3.3, we summarize some results concerning violation of Bell’s inequality in AQFT. Finally, in Section 3.4 we ask whether a local algebra $A(O)$ can be isolated from the influences of its environment.

3.1 Independence of C^* and von Neumann algebras

We first consider notions of independence between a general pair of von Neumann or C^* -algebras.

DEFINITION 83. If e, f are projection operators on a Hilbert space \mathcal{H} , then we let $e \wedge f$ denote the projection onto the closed subspace $e(\mathcal{H}) \cap f(\mathcal{H})$.

FACT 84. Let R be a von Neumann algebra acting on \mathcal{H} . If $e, f \in R$ then $e \wedge f \in R$.

DEFINITION 85 (Schlieder Property). Let R_1, R_2 be von Neumann algebras acting on the Hilbert space \mathcal{H} . We say that the pair (R_1, R_2) satisfies the *Schlieder property* just in case if $e \in R_1$ and $f \in R_2$ are nonzero projections, then $e \wedge f \neq 0$.

The Schlieder property entails that for $e \in R_1, f \in R_2$, if $e, f \neq 0$ and $e, f \neq \mathbb{1}$ then:

$$e \wedge f \neq 0, \quad \neg e \wedge \neg f \neq 0, \quad e \wedge \neg f \neq 0, \quad \neg e \wedge f \neq 0,$$

where $\neg x = \mathbb{1} - x$ is the projection onto the orthogonal complement of $x(\mathcal{H})$. Hence if “ \wedge ” is the analogue of conjunction in classical logic, then the Schlieder property is the analogue of logical independence.

DEFINITION 86. If A, B are C^* -subalgebras of some C^* -algebra C , we let $A \vee B$ denote the C^* -algebra generated by $A \cup B$.

DEFINITION 87 (C^* -Independence). Let A, B be C^* -algebras. We say that the pair (A, B) is *C^* -independent* just in case for any state ω_1 of A and any state ω_2 of B , there is a state ω of $A \vee B$ such that $\omega|_A = \omega_1$ and $\omega|_B = \omega_2$. In other words, each state of A is compatible with each state of B .

The C^* -independence assumption has an obvious operationalist motivation: if Alice is an observer at O_1 and Bob is an observer at O_2 , then C^* -independence amounts to the claim that Alice’s choice to prepare a state cannot in any way obstruct Bob’s ability to prepare a state. Indeed, [Summers and Buchholz, 2005] claim that a failure of C^* -independence could be detected by local observers. On the other hand, C^* -independence could also be regarded as an explication of the notion of the independence of objects:

Two objects A, B are truly independent just in case any state of A is compatible with any state of B ; i.e. there are no logical relations between predications of states to A and B .

Unfortunately, C^* -independence does not imply microcausality.

EXAMPLE 88. We show that C^* -independence does not entail microcausality. (Compare with [Napiórkowski, 1972].) Consider the finite dimensional $*$ -algebra $C(\mathbb{Z}_4) \oplus M_2$, where $C(\mathbb{Z}_4)$ is the abelian $*$ -algebra of dimension 4, and M_2 is the 2×2 matrices over \mathbb{C} . The projection lattice of $C(\mathbb{Z}_4)$ is the Boolean algebra with two atoms; hence it contains logically independent elements e_1, e_2 . Now choose two projections $f_1, f_2 \in M_2$ such that $[f_1, f_2] \neq 0$, and let R_i be the abelian $*$ -subalgebra of $C(\mathbb{Z}_4) \oplus M_2$ generated by the projection $e_i \oplus f_i$.

To see that (R_1, R_2) is C^* -independent, let ω_i be states on the R_i , and let $\lambda_i = \omega_i(e_i \oplus f_i)$. By the logical independence of e_1, e_2 , there is a state ρ of $C(\mathbb{Z}_4)$ such that $\rho(e_i) = \lambda_i$. Then the state $\rho \oplus 0$ on $C(\mathbb{Z}_4) \oplus M_2$ is a common extension of the ω_i since

$$(\rho \oplus 0)(e_i + f_i) = \rho(e_i) = \lambda_i,$$

and a state's value on $e_i \oplus f_i$ determines its value on R_i . Therefore, (R_1, R_2) is C^* -independent. On the other hand, $[e_1 + f_1, e_2 + f_2] = [f_1, f_2] \neq 0$, whence (R_1, R_2) does not satisfy microcausality.

In the previous example, the algebras R_1 and R_2 share a common superselection sector: each commutes with the projection $p = \mathbb{1} \oplus 0$. However, the reduced algebras $pR_i p$ are not C^* -independent. In fact, the diagnosis of this example can be generalized into the following result.

PROPOSITION 89. *Let R_1 and R_2 be von Neumann algebras acting on a Hilbert space \mathcal{H} . If for every projection $e \in Z(R_1 \vee R_2)$, the pair $(eR_1 e, eR_2 e)$ is C^* -independent, then $[R_1, R_2] = \{0\}$.*

Proof. See [Summers and Buchholz, 2005]. ■

DEFINITION 90 (Split Property). Let R_1 and R_2 be von Neumann algebras on \mathcal{H} such that $R_1 \subseteq R'_2$. Then the pair (R_1, R_2) is said to satisfy the *split property* just in case there is a type I factor M such that $R_1 \subseteq M \subseteq R'_2$.


REMARK 91. (i): It is clear that the previous definition is equivalent to saying that (R_1, R'_2) is a ‘split inclusion’ as per Definition 60.

(ii): If (R_1, R_2) satisfies the split property, then under some fairly standard conditions (e.g. R_1 or R_2 is type III), there is a natural $*$ -isomorphism α between $\overline{R_1 \vee R_2}$ and the von Neumann algebra tensor product $R_1 \overline{\otimes} R_2$; by saying that α is ‘natural’, we mean that it extends the map $AB \mapsto A \otimes B$. Furthermore, the $*$ -isomorphism α is spatial, i.e. there is a unitary operator u such that $\alpha(x) = u x u^*$. See [Summers, 1990, p. 212].

(iii): On the other hand, suppose that R is a factor, so that $R \cup R'$ generates $B(\mathcal{H})$ as a von Neumann algebra, i.e. $\overline{R \vee R'} = B(\mathcal{H})$. Then R' is of the same type (I, II, or III) as R [Kadison and Ringrose, 1997, Thm. 9.1.3], and so the von Neumann algebra tensor product $R \overline{\otimes} R'$ is of the same type as R [Kadison and Ringrose, 1997, p. 830]. So if R is type II or III, then $\overline{R \vee R'}$ is strictly larger than, and not isomorphic to $R \overline{\otimes} R'$.

DEFINITION 92 (W^* -Independence). Let R_1 and R_2 be von Neumann algebras acting on \mathcal{H} . The pair (R_1, R_2) is said to be W^* -independent just in case for every normal state φ_1 of R_1 and for every normal state φ_2 of R_2 , there is a normal state φ of $R_1 \vee R_2$ such that $\varphi|_{R_i} = \varphi_i$.

With the assumption of the mutual commutativity of R_1 and R_2 (i.e. microcausality), we have the following implications (see [Summers, 1990, p. 222]):



3.2 Independence of local algebras

We now consider which independence properties hold between pairs of algebras associated with spacelike separated regions. In general, not much can be said about the independence of such algebras. In order to get such results off the ground, we need a stronger notion of spacelike separation.

DEFINITION 93. Two double cones O_1, O_2 are said to be *strictly spacelike separated* just in case there is a neighborhood N of zero such that $O_1 + x$ is spacelike separated from O_2 for all $x \in N$.

PROPOSITION 94. *Suppose that the net $O \mapsto R(O)$ satisfies microcausality, weak additivity, and the spectrum condition. If O_1 and O_2 are strictly spacelike separated, then $(R(O_1), R(O_2))$ satisfies the Schlieder property.*

Proof. See [Schlieder, 1969]. ■

In terms of logical strength, the following concept lies between spacelike separation and strict spacelike separation; furthermore, this concept makes sense for spacetimes without a translation group.

DEFINITION 95. Two double cones O_1 and O_2 are said to be *strongly spacelike separated* just in case there are double cones \tilde{O}_i such that $\bar{O}_i \subseteq \tilde{O}_i$, and \tilde{O}_1, \tilde{O}_2 are spacelike.

FACT 96. If O_1 and O_2 are strictly spacelike separated, then they are strongly spacelike separated.

Of course, the assumptions of Proposition 94 (microcausality, additivity, spectrum) are precisely what is used to derive property B for the net (Proposition 69). So, it is perhaps illustrative to give a simple derivation of the Schlieder property from property B. (Such a result also applies in contexts — e.g. QFT on curved spacetime — where the spectrum condition does not make sense.)

PROPOSITION 97. *Suppose that the net $O \mapsto R(O)$ of von Neumann algebras satisfies microcausality and property B. If O_1 and O_2 are strongly spacelike separated, then $(R(O_1), R(O_2))$ satisfies the Schlieder property.*

Proof. Let O_1 and O_2 be strongly spacelike separated, and let $e_i \in \mathfrak{K}(O_i)$ be projections. Then there are regions \tilde{O}_i such that $\bar{O}_i \subseteq \tilde{O}_i$, and \tilde{O}_1 is spacelike to \tilde{O}_2 . By property B, there are isometries $v_i \in R(\tilde{O}_i)$ such that $v_i v_i^* = e_i$.

Furthermore, $[v_1, v_2] = 0$ and hence $e_1 e_2 = v_1 v_2 (v_1 v_2)^*$. But $v_1 v_2$ is an isometry, and so $v_1 v_2 (v_1 v_2)^* \neq 0$. ■

The split property clearly does not hold for $(R(W), R(W'))$ where W is a wedge region and W' is its causal complement. Indeed, since $R(W)$ and $R(W')$ are type III₁ factors, there can be no *-isomorphism between $R(W) \overline{\otimes} R(W')$ and $R(W) \vee R(W)' = B(H)$. However, if the funnel property holds for the net $O \mapsto R(O)$, then $(R(O_1), R(O_2))$ satisfies the split property when O_1 and O_2 are strictly spacelike separated double cones.

3.3 Bell correlation between von Neumann algebras

We first define a generalized notion of Bell type measurements for a pair of von Neumann algebras.

DEFINITION 98. Let A and B be mutually commuting C^* -subalgebras of some C^* -algebra C . Then we set

$$\mathbb{B}(A, B) \equiv \left\{ (1/2)[a_1(b_1 + b_2) + a_2(b_1 - b_2)] : a_i = a_i^* \in A, b_i = b_i^* \in B, -\mathbb{1} \leq a_i, b_i \leq \mathbb{1} \right\}.$$

Elements of $\mathbb{B}(A, B)$ are called *Bell operators* for (A, B) .

Let r be a Bell operator for (A, B) . It can be shown that $|\varphi(r)| \leq \sqrt{2}$ for each state φ on C [Summers and Werner, 1987]. It is also straightforward to check that if φ is a separable state (i.e. a mixture of product states) then $|\varphi(r)| \leq 1$. Indeed, the Bell measurement correlations in the state φ can be reproduced by a local hidden variable model iff $|\varphi(r)| \leq 1$ [Summers and Werner, 1987; Baez, 1987].

DEFINITION 99. Define the Bell correlation coefficient of a state φ of $A \vee B$ by

$$\beta(\varphi, A, B) = \sup\{ |\varphi(r)| : r \in \mathbb{B}(A, B) \}.$$

If $|\beta(\varphi, A, B)| > 1$, then φ is said to *violate a Bell inequality*, or to be *Bell correlated*.

It is a straightforward exercise to show that if R_1 is an abelian von Neumann algebra and $R_1 \subseteq R'_2$, then for any state φ , $\beta(\varphi, R_1, R_2) \leq 1$. For a sort of converse, Landau [1987] shows that if R_1 and R_2 are nonabelian von Neumann algebras such that $R_1 \subseteq R'_2$, and if (R_1, R_2) satisfies the Schlieder property, then there is *some* state φ that violates Bell's inequality maximally relative to (R_1, R_2) . Similarly, Bacciagaluppi [1994] shows that if A and B are C^* -algebras, then some state violates a Bell inequality for $A \otimes B$ iff both A and B are nonabelian.

When A and B have further properties, we can derive even stronger results. For present purposes, we will simply apply a couple of the known results to the case of AQFT. (See [Summers, 1990] for many more details.)

PROPOSITION 100. *Let R be a type III₁ factor acting on a separable Hilbert space \mathcal{H} . Then every normal state φ of $B(\mathcal{H})$ is maximally Bell correlated across (R, R') , that is $\beta(\varphi, R, R') = \sqrt{2}$.*

Proof. See [Summers and Werner, 1988; Summers and Werner, 1995]. ■

NOTE 101. Prop. 77 tells us that under quite generic conditions, the wedge algebra $R(W)$ is a type III₁ factor. In this case, Prop. 100 tells us that the vacuum is maximally Bell correlated across $(R(W), R(W)')$.

PROPOSITION 102. *Suppose that R_1 and R_2 are von Neumann algebras on \mathcal{H} such that $R_1 \subseteq R_2'$, and (R_1, R_2) satisfies the Schlieder property. If R_1 and R_2 are properly infinite, then there is a dense set of vectors in \mathcal{H} that induce Bell correlated states across (R_1, R_2) .*

Proof. See [Halvorson and Clifton, 2000]. ■

NOTE 103. If a net $O \mapsto R(O)$ of von Neumann algebras on \mathcal{H} satisfies property B and nontriviality, then the hypotheses of Prop. 102 apply to algebras $R(O_1)$ and $R(O_2)$ when O_1 and O_2 are strongly spacelike separated.

Notes: For a comprehensive review of pre-1990 results on independence of local algebras in AQFT, see [Summers, 1990]. For some more recent results, see [Summers, 1997; Florig and Summers, 1997; Rédei, 1998; Halvorson and Clifton, 2000; Summers and Buchholz, 2005].

3.4 Intrinsically entangled states

According to Clifton and Halvorson [2001b], the type III property of local algebras in AQFT shows that it is impossible to disentangle local systems from their environment. To see the argument, recall that it is a standard (perhaps somewhat justified) assumption that the general form of a dynamical evolution T of observables, represented by self-adjoint elements of a C^* -algebra A is given by a completely positive (CP) linear mapping T of A such that $T(\mathbb{1}) = \mathbb{1}$. (Such an assumption is certainly commonplace in, say, quantum information theory.) Here we recall the pertinent definition.

DEFINITION 104. Let A be a C^* -algebra. A linear map T of A is said to be *positive* if $T(a^*a) \geq 0$ for each $a \in A$. T is said to be *completely positive* if for each $n \in \mathbb{N}$, the map $T \otimes \text{id}_n : A \otimes M_n \rightarrow A \otimes M_n$ defined on elementary tensors by

$$(T \otimes \text{id}_n)(a \otimes b) = T(a) \otimes b,$$

is positive. Here M_n is the C^* -algebra of $n \times n$ matrices over \mathbb{C} .

NOTE 105. If $T : \mathfrak{A} \rightarrow \mathfrak{A}$ is positive and $T(\mathbb{1}) = \mathbb{1}$, then for each state ω of \mathfrak{A} , we define $T^*(\omega)$ by $T^*(\omega)(a) = \omega(T(a))$. It follows that T^* is an affine mapping of the state space into itself.

For type I factors, Kraus' theorem [Kraus, 1983] shows that CP maps are “inner.”

THEOREM 106 (Kraus Representation). *If R is a type I_n factor then the following are equivalent for a linear map $T : R \rightarrow R$.*

1. T is completely positive and $T(\mathbb{1}) = \mathbb{1}$.
2. T is the restriction of an automorphism $x \mapsto xux^*$ on an algebra of the form $R \otimes B(H)$.
3. There are positive operators $a_1, \dots, a_n \in R$ such that $\sum_{i=1}^n a_i = \mathbb{1}$ and

$$(1) \quad T(x) = \sum_{i=1}^n a_i^{1/2} x a_i^{1/2}.$$

One special case of Eqn. (1) is the Lüders rule with projection operators e and $\mathbb{1} - e$:

$$T_e(x) = exe + (\mathbb{1} - e)x(\mathbb{1} - e).$$

Furthermore, if the algebra R is type I, we can choose $e \in R$ to be an abelian projection. We have the following result:

If the local algebra R is a type I factor, then there is a universal disentangling operation T_e . That is, no matter what the initial state, the outcome of applying T_e is that the final state is separable.

However, suppose that R has no abelian projections (e.g. R is type III). Then for each nonzero projection $e \in R$, the algebras eRe and $eR'e$ are nonabelian, and hence there is some entangled state φ for the pair $(eRe, eR'e)$. This entangled state is the image under the operation $(T_e)^*$ of some state on $\overline{R \vee R'}$. Hence, the operation T_e does not disentangle all states.

This heuristic argument can be tightened up into a “proof” that no operation on R can disentangle the states of $\overline{R \vee R'}$. See [Clifton and Halvorson, 2001b] for details.

NOTE 107. (i): The Kraus representation theorem is not valid as it stands for type III algebras. Indeed, the Kraus representation theorem is a special case of the Stinespring decomposition theorem [Stinespring, 1955].

(ii): A CP operation on a von Neumann algebra is typically also assumed to be ultraweakly continuous. The continuity of T might be justified on the grounds that it is necessary if T^* is to map normal states to normal states. For objections to the continuity requirement, see [Srinivas, 1980].

4 PROSPECTS FOR PARTICLES

The main application of relativistic QFT is to fundamental *particle* physics. But it is not completely clear that fundamental particle physics is really about particles.

Indeed, despite initial signs that QFT permits a particle interpretation (via Fock space), there are many negative signs concerning the possibility of particle ontology of relativistic QFT. This section is devoted to assessing the status of particles from the point of view of AQFT.

4.1 Particles from Fock space

We begin our investigation of particles with the “story from mother’s knee” about how to give QFT a particle interpretation. (See [Teller, 1995] for one philosopher’s interpretation of this story.) The story begins with a special Hilbert space, called *Fock space*. Now Fock space is just another separable infinite dimensional Hilbert space (and so isomorphic to all its separable infinite dimensional brothers). But the key is writing it down in a fashion that *suggests* a particle interpretation. In particular, suppose that H is the one-particle Hilbert space, i.e. the state space for a single particle. Now depending on whether our particle is a Boson or a Fermion, the state space of a pair of these particles is either $E_s(H \otimes H)$ or $E_a(H \otimes H)$, where E_s is the projection onto the vectors invariant under the permutation $\Sigma_{H,H}$ on $H \otimes H$, and E_a is the projection onto vectors that change signs under $\Sigma_{H,H}$. For present purposes, we ignore these differences, and simply use $H \otimes H$ to denote one possibility or the other. Now, proceeding down the line, for n particles, we have the Hilbert space $H^n \equiv H \otimes \cdots \otimes H$, etc..

A state in H^n is definitely a state of n particles. To get disjunctive states, we make use of the direct sum operation “ \oplus ” on Hilbert spaces. So we define the Fock space $\mathcal{F}(H)$ over H as the infinite direct sum:

$$\mathcal{F}(H) = \mathbb{C} \oplus H \oplus (H \otimes H) \oplus (H \otimes H \otimes H) \oplus \cdots .$$

So, the state vectors in Fock space include a state where there are no particles (the vector lies in the first summand), a state where there is one particle, a state where there are two particles, etc.. Furthermore, there are states that are superpositions of different numbers of particles.

One can spend time worrying about what it means to say that particle numbers can be superposed. But that is the “half empty cup” point of view. From the “half full cup” point of view, it makes sense to count particles. Indeed, the positive (unbounded) operator

$$N = 0 \oplus 1 \oplus 2 \oplus 3 \oplus 4 \oplus \cdots ,$$

is the formal element of our model that permits us to talk about the number of particles.

NOTE 108. In the category of Hilbert spaces, all separable Hilbert spaces are isomorphic — there is no difference between Fock space and the single particle space. If we are not careful, we could become confused about the bearer of the name “Fock space.”

The confusion goes away when we move to the appropriate category. According to Wigner’s analysis [Wigner, 1939], a particle corresponds to an irreducible unitary representation of the identity component \mathcal{P} of the Poincaré group. Then the single particle space and Fock space are distinct objects in the category of representations of \mathcal{P} . The underlying Hilbert spaces of the two representations are both separable (and hence isomorphic as Hilbert spaces); but the two representations are most certainly not equivalent (one is irreducible, the other reducible).

4.2 Fock space from the algebra of observables

The Fock space story is not completely abandoned within the algebraic approach to QFT. In fact, when conditions are good, Fock space emerges as the GNS Hilbert space for some privileged vacuum state of the algebra of observables. We briefly describe how this emergence occurs before proceeding to raise some problems for the naive Fock space story. (We look here only at the symmetric — Bosonic — case. A similar treatment applies to the antisymmetric — Fermionic — case.)

The algebraic reconstruction of Fock space arises from the algebraic version of canonical quantization. Suppose that S is a real vector space (equipped with some suitable topology), and that σ is a symplectic form on S . So, S represents a classical phase space (see Butterfield, this volume). The *Weyl algebra* $\mathfrak{A}[S, \sigma]$ is a specific C^* -algebra generated by elements of the form $W(f)$, with $f \in S$ and satisfying the canonical commutation relations in the Weyl-Segal form:

$$W(f)W(g) = e^{-i\sigma(f,g)/2}W(f+g).$$

Suppose that there is also some notion of spacetime localization for elements of S , i.e. a mapping $O \mapsto S(O)$ from double cones in Minkowski spacetime to subspaces of S . Then, if certain constraints are satisfied, the pair of mappings

$$O \mapsto S(O) \mapsto \mathfrak{A}(O) \equiv C^*\{W(f) : f \in S(O)\},$$

can be composed to give a net of C^* -algebras over Minkowski spacetime. (Here C^*X is the C^* -algebra generated by the set X .)

Now if we are given some dynamics on S , then we can — again, if certain criteria are satisfied — define a corresponding dynamical automorphism group α_t on $\mathfrak{A}[S, \sigma]$. There is then a unique dynamically stable pure state ω_0 of $\mathfrak{A}[S, \sigma]$, and we consider the GNS representation (\mathcal{H}, π) of $\mathfrak{A}[S, \sigma]$ induced by ω_0 . To our delight, we find that the infinitesimal generators $\Phi(f)$ of the one-parameter groups $\{\pi(W(f))\}_{t \in \mathbb{R}}$ behave just like the field operators in the old-fashioned Fock space approach. Furthermore (now speaking non-rigorously), if we define operators

$$\begin{aligned} a(f) &= 2^{-1/2} (\Phi(f) + i\Phi(Jf)), \\ a^*(f) &= 2^{-1/2} (\Phi(f) - i\Phi(Jf)), \end{aligned}$$

we find that they behave like creation and annihilation operators of particles. (Here J is the unique “complex structure” on S that is compatible with the dynamics.)

In particular, by applying them to the vacuum state Ω , we get the entire GNS Hilbert space \mathcal{H} . Finally, if we take an orthonormal basis $\{f_i\}$ of S , then the sum

$$\sum_{i=1}^{\infty} a^*(f_i)a(f_i),$$

is the number operator N . Thus, the traditional Fock space formalism emerges as one special case of the GNS representation of a state of the Weyl algebra.

NOTE 109. The Minkowski vacuum representation (\mathcal{H}_0, π_0) of \mathfrak{A} is Poincaré covariant, i.e. the action $\alpha_{(a, \Lambda)}$ of the Poincaré group by automorphisms on \mathfrak{A} is implemented by unitary operators $U(a, \Lambda)$ on \mathcal{H} . When we say that \mathcal{H} is isomorphic to Fock space $\mathcal{F}(H)$, we do not mean the trivial fact that \mathcal{H} and $\mathcal{F}(H)$ have the same dimension. Rather, we mean that the unitary representation (\mathcal{H}, U) of the Poincaré group is a Fock representation.

Notes: See [Bratteli and Robinson, 1997, Section 5.2] for a detailed account of the reconstruction of Fock space from the Weyl algebra. See also [Clifton and Halvorson, 2001a] and [Halvorson, 2001] for shorter expositions.

4.3 Nonuniqueness of particle interpretations

If we have a representation (\mathcal{H}, π) of the quasilocal algebra \mathfrak{A} such that \mathcal{H} is isomorphic to Fock space, then we can make sense of talk about particles. Furthermore, such representations exist, e.g., the GNS representation of the Minkowski vacuum state ω_0 of the free Bose field. So, in the most simple cases (e.g. free fields on flat spacetime), there is no problem concerning the existence of particle interpretations of the theory.

But there is a problem about *uniqueness*: there are unitarily inequivalent representations of \mathfrak{A} , each of which is isomorphic to Fock space. Furthermore, a result from [Chaiken, 1967; Chaiken, 1968] shows that two inequivalent Fock representations correspond to two number operators that cannot be thought of as notational variants of the same description of reality. Indeed, there are no states of \mathfrak{A} that assign sharp values to both number operators. Hence, the particle interpretations provided by the two Fock representations are mutually exclusive.

The issue of inequivalent Fock representations is treated in depth in [Clifton and Halvorson, 2001a]. For present purposes, we simply note that this worry about nonuniqueness is tied in to a more general worry about inequivalent representations of the quasilocal C^* -algebra \mathfrak{A} . But this more general issue cannot be resolved without reference to recent developments in the theory of superselection sectors (see Sections 7.2 and following). We return to this question in Section 7.

4.4 Problems for localized particles

Suppose that we have settled the uniqueness problem that is raised in the previous subsection — e.g. we have found a good reason for preferring a particular Fock

representation (\mathcal{H}, π) of \mathfrak{A} , and so we have a preferred global number operator N on \mathcal{H} . The next question is whether relativistic QFT is consistent with an ontology of *localized* particles — that is, whether it makes sense to talk about the number of particles in a bounded region O of space.

As pointed out in Section 2.3, the Reeh-Schlieder (RS) Theorem entails that the local algebras of AQFT do not contain operators that annihilate the vacuum. Hence if a number operator has the vacuum as an eigenstate, then there are no local number operators. That is perhaps enough to convince most readers that localized particles are not possible in relativistic QFT. Nonetheless, there have been attempts to bypass the RS Theorem, most notably the proposal of Newton and Wigner (recently resurrected in [Fleming, 2000]). It has been argued that such attempts are not promising [Halvorson, 2001]. Furthermore, it can be shown independently of the full framework of AQFT, and without the RS Theorem, that a positive energy condition combined with microcausality rules out local number operators [Halvorson and Clifton, 2002].

Despite the various No Go results for localized particles in relativistic QFT, the interpretation of experiments in high energy physics seems to require a notion of something causing clicks in detectors, and that a “detector” is fairly well localized in some bounded region of spacetime. A detector corresponds to a positive operator C in \mathfrak{A} , and is “completely reliable” only if it registers 0 identically in the vacuum state, i.e. $C\Omega = 0$. Hence the Reeh-Schlieder Theorem entails that C is not contained in any local algebra. Nonetheless, a notion of approximate localization of C can be salvaged: choose some $A \in \mathfrak{A}(O)$ with $0 \leq A \leq I$, and set

$$C = \int f(x)\alpha_x(A)dx,$$

where f is a smooth function whose Fourier transform has support in the complement of the forward light cone. (The function f automatically has unbounded support.) Then $C\Omega = 0$, and the function f can also be chosen so that C is “close” in the norm topology to an operator in $\mathfrak{A}(O)$.

The notion of approximately localized detectors is employed extensively in Haag-Ruelle scattering theory and recent developments thereof, to which we now turn.

4.5 Particle interpretations generalized: Scattering theory and beyond

It is not true that a representation (\mathcal{K}, π) of \mathfrak{A} must be a Fock representation in order for states in the Hilbert space \mathcal{K} to have an interpretation as particle states. Indeed, one of the central tasks of “scattering theory,” is to provide criteria — in the absence of full Fock space structure — for defining particle states. These criteria are needed in order to describe scattering experiments which cannot be described in a Fock representation, but which need particle states to describe the input and output states.

Haag and Swieca [1965] propose to pick out the n -particle states by means of localized detectors; we call this the *detector criterion*:

A state with at least n -particles is a state that would trigger n detectors that are far separated in space.

Philosophers might worry that the detector criterion is too operationalist. Indeed, some might claim that detectors themselves are made out of particles, and so *defining* a particle in terms of a detector would be viciously circular.

If we were trying to give an analysis of the concept of a particle, then we would need to address such worries. However, scattering theory does not end with the detector criterion. Indeed, the goal is to tie the detector criterion back to some other more intrinsic definition of particle states. The traditional intrinsic definition of particle states is in terms of Wigner’s symmetry criterion:

A state of n particles (of spins s_i and masses m_i) is a state in the tensor product of the corresponding representations of the Poincaré group.

Thus, scattering theory — as originally conceived — needs to show that the states satisfying the detector criterion correspond to an appropriate representation of the Poincaré group. In particular, the goal is to show that there are isometries $\Omega^{\text{in}}, \Omega^{\text{out}}$ that embed Fock space $\mathcal{F}(H)$ into \mathcal{K} , and that intertwine the given representations of the Poincaré group on $\mathcal{F}(H)$ and \mathcal{K} .

Based on these ideas, detailed models have been worked out for the case where there is a mass gap. Unfortunately, as of yet, there is no model in which $\mathcal{H}^{\text{in}} = \mathcal{H}^{\text{out}}$, which is a necessary condition for the theory to have an S-matrix, and to define transition probabilities between incoming and outgoing states. (Here \mathcal{H}^{in} is the image of Fock space in \mathcal{K} under the isometry Ω^{in} , and similarly for \mathcal{H}^{out} .)

Recently, Buchholz and collaborators have claimed that Wigner’s symmetry criterion is too stringent — i.e. there is a more general definition of particle states. They claim that it is only by means of this more general criterion that we can solve the “infraparticles” problem, where massive particles carry a cloud of photons (see [Buchholz *et al.*, 1991]).

Note: For a review of progress in scattering theory in AQFT, see [Haag, 1996, Chapter 6] and [Buchholz and Summers, ND].

5 THE PROBLEM OF VALUE-DEFINITENESS IN AQFT

The “measurement problem” of nonrelativistic QM shows that the standard approach to the theory is impaled on the horns of a dilemma: either (i) one must make ad hoc adjustments to the dynamics (“collapse”) when needed to explain the results of measurements, or (ii) measurements do not, contrary to appearances, have outcomes (see Dickson, this volume, Section 5).

There are two main responses to the dilemma: On the one hand, some suggest that we abandon the unitary dynamics of QM in favor of stochastic dynamics that

accurately predicts our experience of measurement outcomes. On the other hand, some suggest that we maintain the unitary dynamics of the quantum state, but that certain quantities (e.g. position of particles) have values even though these values are not specified by the quantum state. (See Dickson, this volume, Section 5.5 for a more nuanced discussion of the possible responses.)

Both approaches — the approach that alters the dynamics, and the approach with additional values — are completely successful as responses to the measurement problem in nonrelativistic QM. But both approaches run into obstacles when it comes to synthesizing quantum mechanics with relativity. In particular, the additional values approach (e.g. the de Broglie–Bohm pilot-wave theory) appears to require a preferred frame of reference to define the dynamics of the additional values (see [Cushing, 1994, pp. 188–191, 196–198], [Holland, 1995], and [Bohm and Hiley, 1995, Chaps. 11 & 12]), and in this case it would fail the test of Lorentz invariance.

The “modal” interpretation of quantum mechanics is similar in spirit to the de Broglie–Bohm theory, but begins from a more abstract perspective on the question of assigning definite values to some observables. (Following [Bell, 1987], we might call these the “beables” of the theory.) Rather than making an intuitively physically motivated choice of the determinate values (e.g. particle positions), the modal interpretation makes the mathematically motivated choice of the spectral decomposition of the quantum state (i.e. the density operator) as determinate. (See [Dieks and Vermaas, 1998; Vermaas, 1999] for reviews of the modal interpretation; see [Clifton, 1995] for motivation.)

Unlike the de Broglie–Bohm theory, it is not obvious that the modal interpretation must violate the spirit or letter of relativistic constraints, e.g. Lorentz invariance [Dickson and Clifton, 1998, p. 9]. So, it seems that there should be some hope of developing a modal interpretation within the framework of AQFT. This is the starting point for Dieks’ [2000] proposal for a modal interpretation of AQFT. Rather than expound Dieks’ original proposal, we move directly to the criticism in [Clifton, 2000], to which we also refer the reader for further elaboration.

5.1 Clifton–Kitajima classification of modal algebras

Clifton’s critique of the modal interpretation of AQFT is based on a remarkable theorem which classifies all possible “modal subalgebras” of a local von Neumann algebra $\mathfrak{R}(O)$ relative to a state ρ . According to Clifton — and the modal interpreters seem to agree on this point — the algebra \mathfrak{D} , $\mathfrak{D} \subseteq \mathfrak{R}(O)$ of definite local observables should satisfy the following constraints relative to a given state ρ of $\mathfrak{R}(O)$:

DEFINITION 110. Let \mathfrak{R} be a von Neumann algebra, and let ρ be a state of \mathfrak{R} . Then a von Neumann subalgebra \mathfrak{D} of \mathfrak{R} is said to be a *modal algebra* for (\mathfrak{R}, ρ) just in case:

1. (*Value definiteness*) The restricted state $\rho|_{\mathfrak{D}}$ is a mixture of dispersion-free

states. (Definition: A state is *dispersion free* iff it assigns each projection operator either 0 or 1.)

2. (*Definability*) \mathfrak{D} is left invariant under all symmetries of \mathfrak{A} that leave the state ρ invariant.
3. (*Maximality*) \mathfrak{D} is maximal, subject to the first two conditions.

The last requirement is imposed simply to rule out trivial counterexamples to uniqueness — e.g. one could always pick the algebra $\mathbb{C}I$ of scalar multiples of the identity. The second requirement is supposed to explicate the idea that \mathfrak{D} is “picked out by” (i.e. is definable in terms of) the state ρ . We have left the notion of a “symmetry” vague (and we will return to this question in the next subsection), but Clifton takes the symmetries to coincide with the $*$ -automorphisms of \mathfrak{A} , and this is needed for the main result (Theorem 114).

To state this result, we need to define the notion of the centralizer of a state. The following proposition establishes the equivalence of two possible definitions of the centralizer.

PROPOSITION 111. *Let \mathfrak{A} be a von Neumann algebra, let ω be a faithful normal state of \mathfrak{A} , and let σ_t^ω be the modular automorphism group of \mathfrak{A} . Then the following two sets are coextensive:*

1. $\{A \in \mathfrak{A} : \sigma_t^\omega(A) = A, \forall t \in \mathbb{R}\}$
2. $\{A \in \mathfrak{A} : \omega(AB) = \omega(BA), \forall B \in \mathfrak{A}\}$

The proof of Prop. 111 depends on the full apparatus of modular theory. We refer the reader to [Takesaki, 2003, Chap. 8] for details.

DEFINITION 112. It is clear that the set defined in the previous proposition is in fact a von Neumann subalgebra of \mathfrak{A} . We call this subalgebra the *centralizer* of ω in \mathfrak{A} , and we denote it by \mathfrak{A}_ω .

EXAMPLE 113. Let $\mathfrak{A} = \mathfrak{B}(\mathcal{H})$, and let ω be a faithful normal state of \mathfrak{A} . Then ω has the form

$$\omega(A) = \text{Tr}(DA), \quad A \in \mathfrak{A},$$

for some density operator $D \in \mathfrak{A}$. Then $\mathfrak{A}_\omega = \{D\}'$, and $Z(\mathfrak{A}_\omega)$ is the abelian von Neumann algebra $\{D\}''$. In particular, if ω is the maximally mixed state of a type I_n factor, then $\mathfrak{A}_\omega = \mathfrak{B}(\mathcal{H})$, and $Z(\mathfrak{A}_\omega) = \mathbb{C}I$.

The Clifton-Kitajima Theorem shows that there is a unique modal algebra for (\mathfrak{A}, ω) , and in the case that the state ω is faithful, it is $Z(\mathfrak{A}_\omega)$, the center of the centralizer of ω .

THEOREM 114 (Clifton-Kitajima). *Let \mathfrak{A} be a von Neumann algebra acting on a Hilbert space \mathcal{H} , and let ω be a normal state of \mathfrak{A} .*

1. *If ω is faithful then $Z(\mathfrak{A}_\omega)$ is the unique modal algebra for (\mathfrak{A}, ω) .*

2. Generally, the unique modal algebra for (\mathfrak{R}, ω) is $\mathfrak{R} \oplus Z(\mathfrak{R}_\omega)E$, where E is the smallest projection in \mathfrak{R} such that $\omega(E) = 1$, and \mathfrak{R} is the algebra of all bounded operators on $(I - E)(\mathcal{H})$.

The result is proven for the case where ω is faithful in [Clifton, 2000], and for the general case in [Kitajima, 2004].

As pointed out by Clifton [2000], Thm. 114 spells trouble for a modal interpretation of AQFT, because there are many cases where the algebra $Z(\mathfrak{R}_\omega)$ is trivial. (See [Ruetsche and Earman, 2005] for further development of this point.)

1. Let W be a wedge region in Minkowski spacetime, and let Ω be the vacuum state. Then there are no fixed points in $\mathfrak{R}(W)$ of the modular automorphism group σ_t^ω (see the proof of Proposition 77, and also [Driessler, 1975]). Hence, $\mathfrak{R}_\omega = \mathbb{C}I$, and $Z(\mathfrak{R}_\omega) = \mathbb{C}I$.
2. In relativistic QFT, local algebras are the type III₁ hyperfinite factor \mathfrak{R} (see Section 2.5). But \mathfrak{R} has a dense set of ergodic states — states with trivial centralizer. For all these states, $Z(\mathfrak{R}_\omega) = \mathbb{C}I$.

Thus, it makes an enormous difference — at least for the feasibility of the modal interpretation — that local algebras are type III₁. For if local algebras were either type I_∞ or III₀, then there would be good news for the modal interpretation.

PROPOSITION 115. *Let \mathfrak{R} be a type I_∞ factor. Then for every normal state ω of \mathfrak{R} , the unique modal algebra \mathfrak{D}_ω is nontrivial.*

Proof. We have $\mathfrak{D}_\omega = Z(\mathfrak{R}_\omega) = \{D\}''$, where D is the density operator, i.e. the positive operator in \mathfrak{R} that implements the state ω via the trace formula. Furthermore, when \mathfrak{R} is type I_∞, D cannot be a multiple of the identity. ■

PROPOSITION 116. *Let \mathfrak{R} be a type III₀ factor. Then for every faithful normal state ω of \mathfrak{R} , the unique modal algebra \mathfrak{D}_ω is nontrivial.*

Proof. Prop. 3.15 in [Takesaki, 2003, p. 402] entails that \mathfrak{D}_ω has no atomic projections, and hence is infinite dimensional. ■

5.2 What is a symmetry in AQFT?

We note here just one problem with application of the Clifton-Kitajima theorem to AQFT: the notion of symmetry invoked might be too liberal for the setting where we have a *net of algebras over spacetime*, as opposed to a single von Neumann algebra. Clifton’s application of the theorem assumes that any automorphism of \mathfrak{R} is a symmetry. However, if $\mathfrak{R} = \mathfrak{R}(O)$ is just one algebra of an entire net $O \mapsto \mathfrak{R}(O)$, then it is not clear that every automorphism of \mathfrak{R} is a symmetry of the relevant system. What we need is a notion of a symmetry of the net $O \mapsto \mathfrak{R}(O)$.

NOTE 117. A partially ordered set \mathcal{K} can be regarded as a category where for $x, y \in \mathcal{K}$, $\text{Hom}(x, y) = \{(x, y)\}$ if $x \leq y$, and otherwise $\text{Hom}(x, y) = \emptyset$. Let \mathbf{C}^*

be the category with C^* -algebras as objects and $*$ -homomorphisms as arrows. On this conception, a net of C^* -algebras on Minkowski spacetime is a functor $\mathfrak{A} : \mathcal{K} \rightarrow \mathbf{C}^*$ where \mathcal{K} is the category of double cones in Minkowski spacetime, ordered by inclusion, and such that $\mathfrak{A}(\text{Hom}(O_1, O_2))$ is an isometry when $\text{Hom}(O_1, O_2)$ is not empty. (For definitions of functors and natural transformations, see p. 866.)

DEFINITION 118. Let \mathcal{K} be a partially ordered set (e.g. regions in some manifold ordered by inclusion). Let $O \mapsto \mathfrak{A}(O)$ and $O \mapsto \mathfrak{B}(O)$ be nets of C^* -algebras over \mathcal{K} . A *net morphism* $\alpha : \mathfrak{A} \rightarrow \mathfrak{B}$ is a natural transformation between the functors. That is, α consists of a collection of morphisms

$$\{\alpha_O : \mathfrak{A}(O) \rightarrow \mathfrak{B}(O) : O \in \mathcal{K}\},$$

that is natural in O . In other words, for each $f \in \text{Hom}(O_1, O_2)$, $\alpha_{O_2} \circ \mathfrak{A}(f) = \mathfrak{B}(f) \circ \alpha_{O_1}$, which just means that the following diagram commutes

$$\begin{array}{ccc} \mathfrak{A}(O_1) & \xrightarrow{\alpha_{O_1}} & \mathfrak{B}(O_1) \\ \mathfrak{A}(f) \downarrow & & \downarrow \mathfrak{B}(f) \\ \mathfrak{A}(O_2) & \xrightarrow{\alpha_{O_2}} & \mathfrak{B}(O_2) \end{array}$$

FACT 119. Net automorphisms correspond to automorphisms of the quasilocal algebra that leave each local subalgebra globally invariant. To state this precisely, let \mathfrak{A} denote the functor from \mathcal{K} into \mathbf{C}^* , and let \mathfrak{B} denote the inductive limit of \mathfrak{A} . We identify $\mathfrak{A}(O)$ with its image in \mathfrak{B} . Then α is a net automorphism of \mathfrak{A} iff there is an automorphism β of \mathfrak{B} such that

$$\beta|_{\mathfrak{A}(O)} = \alpha_O.$$

Now, given a net \mathfrak{A} with inductive limit \mathfrak{B} , what should we consider as a symmetry of \mathfrak{B} ?

PROPOSAL 120. A symmetry of the net \mathfrak{A} corresponds to a net automorphism α ; i.e. a natural transformation of \mathfrak{A} . That is, a symmetry of \mathfrak{A} corresponds to an automorphism of the quasilocal algebra that leaves each local subalgebra globally invariant.

This first proposal is surely too strict, because it excludes the case of symmetries induced by underlying symmetries of the spacetime. But if \mathcal{K} consists of an appropriate set of regions of spacetime M (i.e. a set that is closed under symmetries of the spacetime), then a symmetry of M will induce an order-preserving bijection F on \mathcal{K} . Note that since F is a functor, $\mathfrak{A} \circ F$ is also a functor. Thus, we consider the following liberalized definition.

PROPOSAL 121. A symmetry of the net \mathfrak{A} consists of a pair (F, α) where F is an order-preserving bijection of \mathcal{K} , and α is a net morphism (natural transformation) from \mathfrak{A} to $\mathfrak{A} \circ F$.

If we accept this proposal, then we must replace Clifton's definability condition with the following modified condition:

Definability-2: Given $O \in \mathcal{K}$, let \mathcal{K}_0 be the full subcategory of \mathcal{K} with objects $\{O_0 : O_0 \leq O\}$, and let \mathfrak{R}_O denote the restriction of the von Neumann algebra valued functor \mathfrak{R} to \mathcal{K}_0 . Then the algebra \mathfrak{D} must be left invariant by all symmetries of \mathfrak{R}_O that preserve the state ρ on $\mathfrak{R}(O)$.

Since not all automorphisms of $\mathfrak{R}(O)$ are symmetries of the net \mathfrak{R}_O , the new definability condition is weaker than the old one: there will typically be more candidates for the role of \mathfrak{D} .

The Clifton-Kitajima Theorem does not apply under the revised definition of symmetries of $\mathfrak{R}(O)$. On the other hand, we are not aware of a positive result showing the existence and uniqueness of subalgebras of $\mathfrak{R}(O)$ that are definite in the state ω and invariant under all net automorphisms that preserve ω . There are suggestive hints such as the result in [Doplicher and Longo, 1984]:

PROPOSITION 122. *Let $(\mathfrak{R}_1 \subseteq \mathfrak{R}_2, \omega)$ be a standard split inclusion of von Neumann algebras. Then there is a unique type I factor \mathfrak{N} such that: (i) $\mathfrak{R}_1 \subseteq \mathfrak{N} \subseteq \mathfrak{R}_2$, and (ii) \mathfrak{N} is invariant under all automorphisms of \mathfrak{R}_2 that preserve both \mathfrak{R}_1 and the state ω .*

Of course, the algebra \mathfrak{N} itself does not have dispersion-free states, and so cannot be the algebra of definite observables. However, the state $\omega|_{\mathfrak{N}}$ is normal, and since \mathfrak{N} is a type I factor, there is a density operator $D \in \mathfrak{N}$ that induces the state in the sense that $\omega(A) = \text{Tr}(DA)$, for all $A \in \mathfrak{N}$. Then assuming that \mathfrak{R}_1 must for some reason be left invariant under symmetries of \mathfrak{R}_2 , the algebra $\mathfrak{D} = \{D\}$ looks like a good candidate for the modal interpreter's set of definite-value observables in \mathfrak{R}_2 in the state ω .

To apply Prop. 122 to AQFT with $\mathfrak{R}_i = \mathfrak{R}(O_i)$, and $O_1 \subseteq O_2$, we would have to assume that the split property holds. Although the split property does not hold in every model, failure of the split property implies a sort of pathology, and it might not be too surprising if there were certain physically pathological cases where the modal interpretation yields a trivial set of definite quantities.

Notes: For recent discussions of adapting the modal interpretation to a relativistic setting, see [Myrvold, 2002; Ruetsche and Earman, 2005].

6 QUANTUM FIELDS AND SPACETIME POINTS

In standard/heuristic presentations of QFT, the fundamental physical quantities (observables, or more generally quantum fields) are operators indexed by space-

time points: $\Phi(x)$ (see t’Hooft, this volume). Based on this fact, at least one philosopher ([Teller, 1995]) describes the ontology of QFT in terms of the idea a *field of operators and their expectation values*. On the other hand, the mathematical approach to QFT (e.g. the Wightman approach) eschews the use of operators at points in favor of operators smeared over space(time) by test-functions: $\Phi(f)$. According to Arnteznius [2003], this fact supports the view that spacetime has no pointlike events, and *a fortiori* that there are no field values at spacetime points.

As QFT became more mathematically rigorous, an intuition developed that it is not only difficult to define the value of a field at a point, but that it is impossible to do so — such quantities simply do not exist. (Compare von Neumann’s critique of Dirac’s delta functions and the notion of pointlike localized particles.) This intuition has sometimes been buttressed by heuristic and operationalist arguments — e.g. Bohr and Petersen’s [1950] argument that it is impossible to measure field strengths at a point. For example, Haag [1996, p. 58] claims that, “a quantum field $\Phi(x)$ at a point cannot be a proper observable.” Even philosophers can be found claiming that, “field operators need to be ‘smeared’ in space” [Huggett, 2000, p. 631, fn. 8].

But the arguments against field operators at a point often confuse questions of measurability with questions of existence, and rarely rise to a level of rigor that is acceptable for drawing metaphysical conclusions. In this section, we review some of the rigorous arguments that exist for and against field quantities at points. We will see that these results do not decisively rule out field quantities at points, but they clarify the interpretive tradeoffs that must be made.

6.1 No Go theorems

In the following three subsections, we review No Go theorems for field operators at spacetime points.

Translation covariance rules out operators at a point

The first no go theorem shows that if there is a continuous unitary representation of the translation group, then for any fixed time t , the field configuration operators $\phi(x, t)$ commute with the field momentum operators $\pi(x', t)$, *even when these operators are associated with the same point*. This result is a serious problem, because $\phi(x, t)$ and $\pi(x, t)$ are supposed to be canonically conjugate (see [Ryder, 1996, p. 131], [Huggett, 1999]):

$$(2) \quad [\phi(x, t), \pi(x', t)] = i\delta(x - x').$$

Moreover, this bad outcome cannot be blamed on any sort of “conflict” between quantum mechanics and relativity, because the bad outcome also holds for non-relativistic theories.

THEOREM 123. *Let $\phi(x, t)$ and $\pi(y, t)$ be fields of operators, either bounded or unbounded and self-adjoint, such that*

$$[\phi(x, t), \pi(y, t)] = 0,$$

when $x \neq y$. (In the unbounded case, we mean that $\phi(x, t)$ and $\pi(y, t)$ are defined on a common dense set \mathcal{D} , and they commute on this set.) If $y \mapsto U(y)$ is a continuous representation of the translation group such that $U(y)\pi(x, t)U(y)^ = \pi(x + y, t)$, for all $x, y \in \mathbb{R}^3$, then*

$$[\phi(x, t), \pi(x, t)] = 0,$$

for all $x \in \mathbb{R}^3$.

Proof. Since this proof only uses field operators on the same time slice, we will suppress reference to t . Suppose first that $\phi(x)$ and $\pi(y)$ are bounded operators. In this case, the mapping:

$$(3) \quad f(y) := [\phi(x), \pi(x + y)] = [\phi(x), U(y)\pi(x)U(y)^*],$$

is a weak-operator continuous function from \mathbb{R}^3 into the bounded operators on \mathcal{H} . Choose a sequence $(y_n)_{n \in \mathbb{N}}$ of nonzero vectors that converges to 0. Since f is continuous, and $f(y_n) = 0$ for all $n \in \mathbb{N}$,

$$(4) \quad [\phi(x), \pi(x)] = f(0) = \lim_{n \rightarrow \infty} f(y_n) = 0.$$

Now suppose that $\phi(x)$ and $\pi(y)$ are unbounded but self-adjoint. Then replace $\pi(x)$ with one of its spectral projections $E_S(x)$, where S is a Borel subset of \mathbb{R} , and replace $\pi(x)$ with one of its spectral projections $F_{S'}(y)$, where S' is a Borel subset of \mathbb{R} . By the preceding argument, $E_S(x)$ and $F_{S'}(y)$ commute. Since this is true for all such pairs of spectral projections, it follows that the spectral projections of $\phi(x)$ commute pairwise with the spectral projections of $\pi(x)$. Hence $\phi(x)$ and $\pi(x)$ are defined on a common dense set \mathcal{D} in \mathcal{H} , and they commute on this dense set. ■

Poincaré covariance rules out operators at a point

For our next two no go theorems, we will need to gather a couple of classic results.

DEFINITION 124. A function $f : \mathbb{R}^n \rightarrow \mathbb{C}$ is said to be *of positive type* just in case for each $c_1, \dots, c_n \in \mathbb{C}$, and each $x_1, \dots, x_n \in \mathbb{R}^n$, we have

$$\sum_{i=1}^n \sum_{j=1}^n \bar{c}_j c_i f(x_i - x_j) \geq 0.$$

THEOREM 125 (Bochner). *Let $f : \mathbb{R}^n \rightarrow \mathbb{C}$ be a continuous function of positive type. Then the Fourier transform of f is a bounded measure on \mathbb{R}^n .*

Proof. For a proof of Bochner's theorem, see [Rudin, 1991, p. 303] and [Folland, 1995, p. 95]. ■

NOTE 126. Only the group structure of \mathbb{R}^n is really needed for Bochner's theorem. So, we are not making any mistake by thinking of Minkowski spacetime as \mathbb{R}^4 in this context.

We will need the following key lemma in all of our subsequent results.

LEMMA 127. *Let f be a continuous positive definite function on \mathbb{R}^n . Then f is the constant 1 function iff the Fourier transform of f is the probability measure with support $\{0\}$.*

The proof of the above lemma is trivial: the Fourier transform of the measure μ with support $\{0\}$ is the function f defined by

$$f(x) = \int_{\mathbb{R}^n} e^{i(x \cdot p)} d\mu(p) = e^{i(x \cdot 0)} = 1.$$

But the Fourier transformation is a bijection between complex Radon measures on \mathbb{R}^n and bounded continuous functions on \mathbb{R}^n .

DEFINITION 128. We say that a measure μ on Minkowski spacetime is *Lorentz invariant* just in case $\mu(\Lambda(S)) = \mu(S)$ for each Borel subset S of M , and each homogeneous Lorentz transformation Λ , where $\Lambda(S) = \{\Lambda(x) : x \in S\}$.

Clearly, the only Lorentz invariant probability measure on Minkowski spacetime is the measure supported on $\{0\}$ (the unique fixed point of the homogeneous Lorentz group). The following result is the "Fourier transformed" version of that fact.

LEMMA 129. *Let M be Minkowski spacetime. If $f : M \rightarrow \mathbb{C}$ is a continuous function of positive type such that $f(\Lambda x) = f(x)$ for each Lorentz transformation Λ , then f is constant.*

Sketch of proof. By Bochner's theorem, if $f : M \rightarrow \mathbb{C}$ is a continuous function of positive type, then f is the Fourier transform of a bounded measure μ on M . It is straightforward to verify that if f is Lorentz invariant then so is μ . But a bounded, Lorentz invariant measure is supported on $\{0\}$. By Lemma 127, the Fourier transform of μ is a constant function. Therefore, $f = 1$ is constant. ■

FACT 130. Let U be a unitary representation of the translation group on a Hilbert space \mathcal{H} . Then the following are equivalent:

1. The spectrum of the representation U is Δ ;
2. For every $u, v \in \mathcal{H}$, the function $f : \mathbb{R}^n \rightarrow \mathbb{C}$ given by

$$f(x) = \langle u, U(x)v \rangle, \quad x \in \mathbb{R}^n,$$

has Fourier transform with support in Δ .

Finally, the following is our core lemma for the next two results.

LEMMA 131. *Let $A : M \rightarrow \mathfrak{B}(\mathcal{H})$ be an operator valued function, and let U be a unitary representation of the translation group on \mathcal{H} such that $U(x)A(0)U(x)^* = A(-x)$ for all $x \in M$. Define a function $f : M \rightarrow \mathbb{C}$ by*

$$f(x) = \langle \Omega, A(x)A(0)\Omega \rangle = \langle \Omega, U(x)^*A(0)U(x)A(0)\Omega \rangle.$$

If f is constant, then there is a $c \in \mathbb{C}$ such that $A(x)\Omega = c\Omega$ for each $x \in M$.

Proof. Let $\psi = A(0)\Omega$. Then $f(x) = f(0)$ is expressed as

$$\langle \psi, U(x)\psi \rangle = \langle \psi, \psi \rangle = \|\psi\|^2.$$

But we also have $\|\psi\| = \|U(x)\psi\|$ since $U(x)$ is unitary. Hence

$$\langle \psi, U(x)\psi \rangle = \|\psi\| \cdot \|U(x)\psi\|,$$

and the Cauchy-Schwartz inequality entails that $U(x)\psi = \psi$ for all x . That is, $U(x)A(0)\Omega = A(0)\Omega$. Note in addition that $U(x)A(y)\Omega = U(x+y)A(0)\Omega = A(0)\Omega$. Hence all vectors $A(x)\Omega$ are invariant under the translation group. ■

Now, the second no go theorem (due to [Wizimirski, 1966]) shows that there is no nontrivial Poincaré covariant field of bounded operators on Minkowski spacetime.

THEOREM 132. *Suppose that $A : M \rightarrow \mathfrak{B}(\mathcal{H})$ is an operator-valued function, and U is a continuous unitary representation of the Poincaré group on \mathcal{H} such that:*

1. $U(y, \Lambda)A(x)U(y, \Lambda)^* = A((\Lambda x) - y)$, for all $(y, \Lambda) \in \mathcal{P}$ and $x \in M$;
2. *There is a unique (up to scalar multiples) translation-invariant vector $\Omega \in \mathcal{H}$.*

Then there is a $c \in \mathbb{C}$ such that $A(x)\Omega = c\Omega$ for all $x \in M$.

NOTE 133. (i): The assumption of the uniqueness of Ω might seem unwarranted. But under some fairly standard conditions, this assumption can be derived. See Section 2.2. (ii): This theorem makes no assumption about commutation relations between operators $A(x)$ and $A(y)$.

Proof [of Theorem 132.] Define a function $f : M \rightarrow \mathbb{C}$ by

$$f(x) = \langle \Omega, A(x)^*A(0)\Omega \rangle, \quad x \in M.$$

By condition 2 we have $U(x)\Omega = \Omega$. Hence by condition 1 we have $A(x)^* = U(x)A(0)^*U(x)^*$, and hence

$$f(x) = \langle A(0)\Omega, U(x)^*A(0)\Omega \rangle,$$

which is obviously positive definite. Furthermore, since $x \mapsto U(x)^*$ is weakly continuous, f is continuous.

Now we establish that $f(\Lambda(x)) = f(x)$ for all $x \in M$ and all Lorentz transformations Λ . We have

$$\begin{aligned} f(\Lambda x) &= \langle \Omega, A(\Lambda x)^* A(0)\Omega \rangle \\ &= \langle \Omega, U(0, \Lambda)A^*(x)U(0, \Lambda)^{-1}A(0)\Omega \rangle \\ &= \langle U(0, \Lambda)^{-1}\Omega, A(x)^*U(0, \Lambda)^{-1}A(0)U(0, \Lambda)\Omega \rangle \\ &= \langle \Omega, A(x)^*A(\Lambda(0))\Omega \rangle \\ &= \langle \Omega, A(x)^*A(0)\Omega \rangle \\ &= f(x). \end{aligned}$$

Thus, Lemma 129 entails that f is constant, and Lemma 131 entails that there is a $c \in \mathbb{C}$ such that $A(x)\Omega = c\Omega$ for all $x \in M$. ■

Microcausality and Spectrum Condition rule out operators at a point

The final no go theorem, originally by Wightman [1964] invokes both microcausality and the spectrum condition. (See [Horuzhy, 1990, p. 46] and [Baumgärtel and Wollenberg, 1992, p. 115] for alternative proofs.)

THEOREM 134. *Suppose that $A : M \rightarrow \mathfrak{B}(\mathcal{H})$ is an operator valued function, and U is a continuous unitary representation of the translation group on \mathcal{H} such that:*

1. $[A(x), A(y)] = 0$ when x and y are spacelike separated;
2. $U(x)A(y)U(x)^* = A(y - x)$, for all $x, y \in M$;
3. U satisfies the spectrum condition.
4. There is a unique translation invariant vector $\Omega \in \mathcal{H}$.

Then there is a $c \in \mathbb{C}$ such that $A(x)\Omega = c\Omega$ for all $x \in M$.

Proof. As above, define $f : M \rightarrow \mathbb{C}$ by

$$f(x) = \langle \Omega, A(x)A(0)\Omega \rangle, \quad x \in M.$$

Fix a nonzero spacelike vector x . Then by condition 1,

$$U(x)^*A(0)U(x)A(0) = A(x)A(0) = A(0)A(x) = A(0)U(x)^*A(0)U(x).$$

Therefore,

$$\begin{aligned} f(x) &= \langle \Omega, U(x)^*A(0)U(x)A(0)\Omega \rangle = \langle \Omega, A(0)U(x)^*A(0)\Omega \rangle \\ &= \langle \Omega, A(0)U(-x)A(0)\Omega \rangle = f(-x). \end{aligned}$$

Now consider the function $F : \mathbb{R} \rightarrow \mathbb{C}$ given by $F(t) = f(tx)$, so that $F(t) = F(-t)$. By condition 3, the Fourier transform of f is supported in the forward

light cone. Hence, the Fourier transform of F is supported in $[0, +\infty)$. But since $F(t) = F(-t)$, the Fourier transform of F is also supported in $(-\infty, 0]$. Therefore, the Fourier transform of F is the point mass at $\{0\}$. By Lemma 129, F is constant. Finally, since any two points in M can be connected by two spacelike vectors, we can apply the previous procedure twice to show that f is constant. Therefore, by Lemma 131, there is a $c \in \mathbb{C}$ such that $A(x)\Omega = c\Omega$ for all $x \in M$. ■

COROLLARY 135. *Let $O \mapsto \mathfrak{R}(O)$ be a net of von Neumann algebras acting irreducibly on a Hilbert space \mathcal{H} , and let U be a strongly continuous unitary representation that implements the action of the translation group on the net \mathfrak{R} . Suppose that the net satisfies microcausality (assumption 43). Suppose that U satisfies the spectrum condition, and that there is a translation invariant vector $\Omega \in \mathcal{H}$. Then for each point $x \in M$,*

$$\bigcap_{\{O \in \mathcal{K} : x \in O\}} \mathfrak{R}(O) = \mathbb{C}I.$$

Proof. Fix $x \in M$, and fix a double cone $x \in O$. Choose an arbitrary operator, denoted by $A(x)$, in

$$\bigcap_{\{O \in \mathcal{K} : x \in O\}} \mathfrak{R}(O).$$

Now for general $y \in M$, define

$$A(y) = U(x - y)A(x)U(x - y)^*,$$

so that the mapping $A : M \rightarrow \mathfrak{B}(\mathcal{H})$ automatically satisfies condition 2 of Theorem 134. Furthermore, since the net \mathfrak{R} satisfies microcausality, and the unitary group U implements the translations on \mathfrak{R} , the mapping A satisfies condition 1 of Theorem 134. It then follows that there is a $c \in \mathbb{C}$ such that $A(x) = cI$. Since x was an arbitrary element of M , the result is proven. ■

6.2 Go theorems

Why should we care if $\Phi(x)$ cannot be taken to denote any non-trivial operator on Hilbert space? Does this have any implications for the interpretation of QFT? After all, for any neighborhood O of x , we can find a test-function f that is supported in O , and we can replace the non-denoting term “ $\Phi(x)$ ” with the denoting term “ $\Phi(f)$ ”. In fact, couldn’t we think of “ $\Phi(x)$ ” as a name for the sequence $\{\Phi(f_n)\}_{n=1}^\infty$, where $\{f_n\}_{n=1}^\infty$ is a sequence of test-functions that converges to the delta-function at x ? More precisely, it seems that we could even attempt to define an expectation value for the pseudo-operator $\Phi(x)$ as follows: If ρ is a state of the quantum field, define:

$$(5) \quad \rho(\Phi(x)) := \lim_{n \rightarrow \infty} \rho(\Phi(f_n)).$$

In this section, we make this idea precise in two Go Theorems for field quantities at points. The first result we report from the work of Rehberg and Wollenberg [Rehberg and Wollenberg, 1986; Wollenberg, 1986] (see also [Fredenhagen and Hertel, 1981], [Bostelmann, 2000; Bostelmann, 2004]). This result shows that within the Wightman framework, a quantum field at a point can be represented by a sesquilinear form. The second result shows that if we drop the requirement of continuity on our representation of the translation group, then quantum fields at points can be represented by self-adjoint operators.

Quantum fields as sesquilinear forms

DEFINITION 136. Let \mathcal{H} be a Hilbert space. A *sesquilinear form* on \mathcal{H} is a linear subspace $D(t)$ of \mathcal{H} and a mapping $t : D(t) \times D(t) \rightarrow \mathbb{C}$ that is antilinear in the first argument, and linear in the second argument. The form t is said to be *densely defined* just in case $D(t)$ is dense in \mathcal{H} . The form t is said to be *symmetric* just in case $t(\varphi, \psi) = \overline{t(\psi, \varphi)}$ for all $\varphi, \psi \in D(t)$. The form t is said to be *positive* just in case $t(\psi, \psi) \geq 0$ for all $\psi \in D(t)$.

DEFINITION 137. If t is a sesquilinear form on \mathcal{H} then we define the associated *quadratic form* by $t(\psi) = t(\psi, \psi)$ for all $\psi \in D(t)$. A positive quadratic form t is said to be *closed* just in case for any sequences $(\psi_n)_{n \in \mathbb{N}}$ in $D(t)$ if $\psi_n \rightarrow \psi$ and $t(\psi_n - \psi_m) \rightarrow 0$, then $\psi \in D(t)$ and $t(\psi_n - \psi) \rightarrow 0$.

NOTE 138. A densely defined, symmetric sesquilinear form is a *prima facie* candidate to represent a physical quantity or an observable. Since t is symmetric, the corresponding quadratic form is real-valued. Hence, for each unit vector $\psi \in D(t)$, we might say that the “expectation value” of t in state ψ is $t(\psi)$. Indeed, at first glance, the expectation value mapping $t \mapsto t(\psi)$ seems to have all the same properties as the corresponding expectation mapping for operators.

THEOREM 139. Let $\Phi(\cdot)$ be a Wightman field on the Hilbert space \mathcal{H} . That is, Φ maps elements of a test-function space $\mathcal{S}(\mathbb{R}^4)$ to unbounded operators on \mathcal{H} with some common dense domain \mathcal{D} . Let $(\delta_n)_{n \in \mathbb{N}}$ be a sequence of test-functions whose support shrinks to the point x . Then for each $u, v \in \mathcal{D}$, the sequence

$$\langle u, \Phi(\delta_1)v \rangle, \langle u, \Phi(\delta_2)v \rangle, \langle u, \Phi(\delta_3)v \rangle, \dots,$$

converges to a finite complex number, which we denote by $\langle u, \Phi(x)v \rangle$. The map $u, v \mapsto \langle u, \Phi(x)v \rangle$ is a sesquilinear form with domain \mathcal{D} , which we denote by $\Phi(x)$.

Proof. See [Baumgärtel and Wollenberg, 1992, p. 332] and [Rehberg and Wollenberg, 1986; Wollenberg, 1986]. ■

NOTE 140. One naturally wishes to have a version of this theorem in a more purely algebraic setting. Such a result might be available in the context of the scaling algebras of [Buchholz and Verch, 1995; Buchholz, 1998].

The result is surprising for a couple of reasons. We might have thought that the reason $\Phi(x)$ is not an operator is because the expectation values $\langle u, \Phi(\delta_n)u \rangle$ grow without bound as the test-functions δ_n shrink to a point — i.e. there is some sort of divergence. But Theorem 139 shows that conjecture to be false. The obstruction to $\Phi(x)$'s becoming an operator must lie somewhere else.

So, we have several No Go Theorems against quantum fields as operators (even unbounded operators), and one Go Theorem for quantum fields as sesquilinear forms. What should we conclude from these apparently conflicting results? Should we say that there is a field quantity at the point x , or not?

To answer this question we need to think harder about the relation between operators on a Hilbert space and physical quantities. Why was it that we thought that physical quantities correspond to operators? If we suppose that an unbounded operator can represent a quantity, then must that operator be self-adjoint (i.e. must A coincide with A^* on a common dense domain), or does it suffice to satisfy some weaker condition? Does any symmetric sesquilinear form have all the features necessary to represent a physical quantity? In order to get clear on these questions, it might help to get clear on the mathematical details of the relationship between sesquilinear forms and operators. Fortunately, there are quite a few results in this direction.

Clearly, every linear (possibly unbounded) operator T on \mathcal{H} defines a sesquilinear form with domain $D(T)$ via the equation

$$(6) \quad t(\psi, \varphi) = \langle \psi, T\varphi \rangle.$$

On the other hand, it is less clear when an arbitrary form t corresponds an operator via Eqn. (6).

DEFINITION 141. A sesquilinear form t on \mathcal{H} is said to be *bounded* just in case there is a $n \in \mathbb{N}$ such that $|t(\varphi, \psi)| \leq n$ whenever $\varphi, \psi \in D(t)$ with $\|\varphi\|, \|\psi\| \leq 1$.

PROPOSITION 142. *There is a one-to-one correspondence between densely defined, bounded sesquilinear forms on \mathcal{H} and elements of $\mathfrak{B}(\mathcal{H})$. In particular, if t is bounded sesquilinear form on \mathcal{H} then there is a unique operator $T \in \mathfrak{B}(\mathcal{H})$ such that $t(\varphi, \psi) = \langle \varphi, T\psi \rangle$ for all $\varphi, \psi \in \mathcal{H}$. Furthermore, t is symmetric iff T is self-adjoint.*

Proof. See [Kadison and Ringrose, 1997, Theorem 2.4.1]. ■

PROPOSITION 143. *If t is a densely defined, positive, closed, quadratic form, then there exists a unique positive operator T on \mathcal{H} such that the domain of $T^{1/2}$ is $D(t)$ and*

$$t(\varphi, \psi) = \langle T^{1/2}\varphi, T^{1/2}\psi \rangle,$$

for all $\varphi, \psi \in D(t)$. In particular, $t(\varphi, \psi) = \langle \varphi, T\psi \rangle$ for all $\varphi, \psi \in D(t)$.

NOTE 144. The previous Proposition is useful in showing when a number operator N can be defined in a representation of the Weyl algebra $\mathfrak{A}[S, \sigma]$. For details, see [Clifton and Halvorson, 2001a] and [Bratteli and Robinson, 1997, p. 27].

The previous two propositions do not apply to the sesquilinear form $\Phi(x)$ because it is neither bounded nor positive. Furthermore, there is no known (to the author) characterization of when a symmetric sesquilinear form admits a representation as an operator — although there are some partial results in this direction (see [McIntosh, 1970]). It is clear that $\Phi(x)$ is not an operator; it is unclear what features operators have that $\Phi(x)$ lacks, and whether these features are necessary for a mathematical object to represent a quantity. Accordingly, it is unclear whether or not $\Phi(x)$ represents an element of reality.

Quantum fields as operators on non-separable Hilbert space

Our second Go result for quantum field operators at a point is really just a sketch of an example. We take a nonseparable Hilbert space H that can represent states of particles with point positions (compare with [Halvorson, 2004]). We then apply the standard second quantization procedure — which does not depend on the one-particle space being separable — to obtain a Fock space $\mathcal{F}(H)$, and self-adjoint field operators $\phi(x), \pi(x)$ indexed by points in \mathbb{R} .

Let $H = l_2(\mathbb{R})$ be the Hilbert space of square-summable sequences over \mathbb{R} ; i.e. an element f of $l_2(\mathbb{R})$ is a mapping from \mathbb{R} into \mathbb{C} such that f vanishes at all but countably many points, and $\sum_{x \in \mathbb{R}} |f(x)|^2 < \infty$. The inner product on $l_2(\mathbb{R})$ is given by

$$(7) \quad \langle f, g \rangle = \sum_{x \in \mathbb{R}} \overline{f(x)}g(x).$$

Let $\mathcal{F}(H)$ be the Fock space over H . For each $x \in \mathbb{R}$, we let $\delta_x \in l_2(\mathbb{R})$ denote the characteristic function of $\{x\}$; the set $\{\delta_x : x \in \mathbb{R}\}$ is an (uncountably infinite) orthonormal basis for $l_2(\mathbb{R})$. For any $x \in \mathbb{R}$, we define the creation operator $a(x)$ by:

$$(8) \quad a(x)(f_1 \otimes \cdots \otimes f_n) := \delta_x \otimes f_1 \otimes \cdots \otimes f_n.$$

As in the standard case, we verify that $a^-(x) + ia^+(x)$ and $a^+(x) - ia^-(x)$ are preclosed, i.e. that the closure of the graphs of these operators are graphs of linear operators (see [Kadison and Ringrose, 1997, p. 155]), which we denote by

$$(9) \quad \phi(x) = \overline{a^-(x) + ia^+(x)},$$

$$(10) \quad \pi(x) = \overline{a^+(x) - ia^-(x)}.$$

It then follows that $\phi(x)$ and $\pi(x)$ are self-adjoint, and on a dense domain \mathcal{D} in $\mathcal{F}(H)$, we have

$$(11) \quad [\pi(x), \phi(x')] = i \langle \delta_x, \delta'_{x'} \rangle = i \delta_0(x - x'),$$

where now δ_0 is a completely legitimate mathematical object — viz. the probability measure supported on $\{0\}$.

Consider the (discontinuous) representation $x \mapsto V(x)$ of the translation group on $l_2(\mathbb{R})$ defined on the basis elements $\{\delta_y : y \in \mathbb{R}\}$ by

$$(12) \quad V(x)\delta_y = \delta_{y-x}.$$

Let Γ be the ‘Fock functor’; i.e. Γ maps a unitary operator V on the single particle space H to the corresponding operator

$$I \oplus V \oplus (V \otimes V) \oplus \dots,$$

on $\mathcal{F}(H)$. Then $x \mapsto U(x) := \Gamma(V(x))$ is a *discontinuous* representation of the translation group on $\mathcal{F}(H)$, and

$$(13) \quad U(x)^*\phi(y)U(x) = \Phi(y-x).$$

Thus, $(\phi(\cdot), \pi(\cdot), \mathcal{F}(H), U)$ is a field system over $l_2(\mathcal{R})$, where $x \mapsto U(x)$ is a discontinuous unitary representation. We could then use the field system to define a net $O \mapsto \mathfrak{A}(O)$ of von Neumann algebras on $\mathcal{F}(H)$. But this net of course fails the spectrum condition, because the representation of the translation group is not continuous.

The model just described is probably too unwieldy to be of much use in describing real physical situations. Furthermore, there is no good reason to think that the procedure we followed generalizes to the case of interacting theories, where pointlike localized operators are needed to make sense of products of field operators. However, we hoped to show that it is conceivable that we can get by without some of the technical assumptions of the No Go Theorems. So, we should think very carefully before we try to use these theorems to draw conclusions about how QFT must be interpreted.

6.3 Field interpretations of QFT

In Section 4, we saw that there are severe obstacles to a particle interpretation of QFT. One might argue then, by a process of elimination, that we should adopt a ‘‘field interpretation’’ of QFT (see e.g. [Teller, 1995; Huggett, 2000]). But if we examine the field interpretation on its own merits, its not clear that it is better off than the particle interpretation.

In constructing canonical free theories (e.g. the free Bose and Fermi fields), one begins with a Hilbert space H which can be interpreted either as a ‘‘single particle space’’ (i.e. space of wavefunctions of a single quantum mechanical particle) or as a space of configurations of a classical field. Corresponding to these two interpretations, there are two ways to construct the Hilbert space of the quantum field theory:

1. Second quantization: the Hilbert space of the quantum field is the Fock space $\mathcal{F}(H)$ over H . (See Section 4.1).

2. Field quantization: the Hilbert space of the quantum field is is space $L_2(H, d)$ of ‘square integrable’ functions from H into \mathbb{C} relative to the isonormal distribution d on H .

(In a rigorous treatment, elements of $L_2(H, d)$ are not really functions. See [Baez *et al.*, 1992, Section 1.3] for details.) The free field theories constructed by these two methods are known to be unitarily equivalent. However, the field quantization approach lends itself more naturally to a field interpretation. Indeed, in a recent survey of the foundations of QFT [Huggett, 2000], one finds the suggestive notation:

$$\Psi(\phi), \quad \phi \in L_2(\mathbb{R}^{3n}),$$

for a function on the space $H := L_2(\mathbb{R}^{3n})$ of wavefunctions. Thus, it seems that a quantum field state can be interpreted as a superposition of classical field configurations in the same sense that a wavefunction of n particles can be interpreted as a superpositions of classical configurations of n particles.

However, there are difficulties with this approach. First, the field operators $\Phi(x)$ on $L_2(H, d)$ are the precise analogues of the position operators Q_i for a system of particles. That is, there is a natural interpretation of a function $\Psi \in L_2(H, d)$ as a probability distribution over the family $\{\Phi(x) : x \in M\}$ of mutually commuting operators. But the No Go theorems for the operators $\Phi(x)$ militate against interpreting Ψ as a probability distribution over classical field configurations. More directly, since d assigns zero measure to points in H (i.e. to individual field configurations), characteristic functions of singleton subsets of H — i.e. determinate field configurations — are identified with the zero vector in $L_2(H, d)$. That is, there is no state of the quantum field in which it is in a definite configuration.

It follows from the preceding considerations that the No Go theorems for fields operators at a point undermine the field interpretation of QFT in the same way that No Go theorems for number operators undermine the particle interpretation. Thus, we should be wary of arguments for field interpretations based on problems with particle interpretations.

6.4 Points of time?

The preceding results were aimed at answering the question of whether there can be field operators at a spacetime point. Suppose that we concede that there cannot be, and we proceed with the standard mathematically rigorous approach to, say, the free Bose field, where field operators are smeared over *space* by test functions (see e.g. [Araki, 1963]). In this case, quantities are not tied to spacetime points, but they are tied to pointlike times. However, some claim that in general, the quantities will also have to be fuzzy in time. For example, according to Haag,

Renormalization theory suggests that it is essential to smear out Φ both in space and time, in contrast to the case of free fields, where an averaging over 3-dimensional space at a fixed time is sufficient. Due to

the stronger singularities, one cannot assume well-defined commutation relations of fields at equal time. [Haag, 1996, p. 59]

But such claims are speculative — we know of no theorems that prove that interacting fields must be smeared out in time. So, at the present time we have no particularly good reason to conclude that time is pointless.

7 THE PROBLEM OF INEQUIVALENT REPRESENTATIONS

The philosophy of local quantum physics (in Haag’s terminology) is that the theoretical parts of QFT (e.g. unobservable fields, gauge group) should not count as part of the given data. Instead, the abstract net \mathfrak{A} of observable algebras should be taken as the primitive. Following the terminology in [Ruetsche, 2002], we define ‘Algebraic Imperialism’ as the position that:

The physical content of a quantum field theory is encoded in the net $O \mapsto \mathfrak{A}(O)$, the subgroup of $\text{Aut}(\mathfrak{A})$ corresponding to physical symmetries (including dynamics), and the states on the quasilocal algebra \mathfrak{A} . A representation (\mathcal{H}, π) of \mathfrak{A} may be an aid to calculation, but has no ontological significance.

Such an attitude might seem incomprehensible to those steeped in the traditional Hilbert space formalism of QM. Indeed, where is the Hamiltonian, where are the transition probabilities, and how do we describe measurements? The very abstractness and generality of the algebraic formalism seems to empty it of a great deal of the content we expect in a physical theory.

However, some of these worries about lack of content of the abstract algebraic formalism are ill founded. Indeed, the GNS theorem (Thm. 17) shows that all the Hilbert spaces we will ever need are hidden inside the algebra itself. Furthermore, much of the vocabulary one learns to speak in elementary QM can be defined within this purely abstract setting. For example, for a definition of transition probabilities between states, see [Roberts and Roepstorff, 1968]; and for a definition of measurement probabilities, see [Wald, 1994].

But it is not true — at least on the face of it — that *all* pieces the traditional vocabulary of QFT can be reproduced in the algebraic setting. For example, the quasilocal algebra does not contain a number operator, and probability distributions over the spectrum of the number operator cannot be defined in terms of expectation values on \mathfrak{A} (see [Clifton and Halvorson, 2001a]). What is perhaps even worse is that by beginning with a net $O \mapsto \mathfrak{A}(O)$ of *observable* algebras, we have effectively closed our eyes to the existence of unobservable fields, which do not generally commute with each other at spacelike separation. Thus, we seem to have no way to account for deep theoretical facts of QFT such as the connection between spin (commutation relations of field operators) and statistics.

Worries such as these might push us towards the second main position on the issue of representations, which Ruetsche [2002] calls Hilbert Space Conservatism:

The theory is not the net $O \mapsto \mathfrak{A}(O)$, but the net plus one specific representation (\mathcal{H}, π) .

In fact, Hilbert Space Conservatism might be thought of as the default view of most workers in mainstream (Lagrangian) QFT, since the abstract algebra (and its representations) do not play a central role there.

But as with many realist views, the Conservative view faces epistemological difficulties: How do we decide which is the correct representation? In this case, the difficulty is particularly severe, because it can be proven mathematically that the predictions of states within any one representation can be reproduced to arbitrarily high accuracy by the states in *any other* representation.² (This is due to the fact that since \mathfrak{A} is simple, Fell's theorem implies that the states in any folium are weak* dense in the state space.)

Nonetheless, it is tempting to think that the algebraic formalism is creating an interpretive problem. That is, it is tempting to think that if we stick to the old-fashioned way of doing QFT, the issue of inequivalent representations does not arise, and so neither does this interpretive dilemma. So, are inequivalent representations telling us something of foundational importance, or are they just mathematical playthings?

The motivating arguments for the algebraic approach have included the existence of inequivalent representations of the canonical commutation relations, as well as physical effects associated with Rindler-Fulling quanta. Nonetheless, these arguments have been resisted for various reasons, e.g. there is a suspicion that the Rindler vacuum representation is inferior, as a description of reality, to the Minkowski vacuum representation. So, in the following sections, we discuss another motivating argument for the algebraic approach — viz. superselection rules. It is in the analysis of superselection rules that the algebraic approach most clearly displays its beauty, utility, and foundational importance.

7.1 Superselection rules

In a now famous paper, Wick, Wightman, and Wigner [1952] argue that there is a physical system with state space H , and state vectors $\psi_1, \psi_2 \in H$ such that the linear combinations

$$(14) \quad 2^{-1/2}(\psi_1 + e^{i\theta}\psi_2), \quad \theta \in [0, 2\pi),$$

give rise to ‘empirically indistinguishable’ states. When this occurs, Wick *et al.* say that there is a ‘superselection rule’ between ψ_1 and ψ_2 ; alternatively, ψ_1 and ψ_2 lie in different ‘superselection sectors.’ We put ‘empirically indistinguishable’ in scare quotes, because the literature has been anything but clear about the nature

²This way of stating the problem is biased, and depends on taking “predictions of a representation” to mean expectation values of observables in the abstract algebra. If we also include expectation values of observables in the weak closure $\pi(\mathfrak{A})^-$, and expectation values of unbounded operators on \mathcal{H} , then the story becomes more complicated. Compare with [Clifton and Halvorson, 2001b].

of the relation between the states in Eqn. (14). Are the states in Eqn. (14) merely *empirically* indistinguishable, or is there a stronger sense in which these two states are equivalent? If the indistinguishability is empirical, how strong is the modal force? Do these states predict the same empirical phenomena in all physically possible worlds, or is their indistinguishability due to some special features (e.g. initial conditions) in our world? In this paper, we will not attempt to resolve these important questions about the nature of superselection rules.³ Rather, we will content ourselves with explaining the proposal of Doplicher, Haag, and Roberts (DHR) for making Wick *et al.*'s notion precise within the context of AQFT.

The first approaches to superselection rules involved an *ad hoc* butchery of the state space and of the algebra of observables, with a resulting confusion about what should count as the states and observables (or quantities) of the resulting theory. One begins with a Hilbert space H with unit vectors giving pure states, and with self-adjoint elements of $B(H)$ as observables (or quantities). One then gives a heuristic argument for the claim that a superselection rule holds between the state vectors in some subspace H_1 and the state vectors in the complementary subspace $H_2 := (H_1)^\perp$. On the basis of this argument, the state space H is reduced to the union of H_1 and H_2 ; that is, a linear combination of a vector in H_1 and H_2 is no longer thought to be a possible (pure state) of the theory; the pure state vectors lie in *either* H_1 *or* H_2 . Equivalently, the algebra of observables $B(H)$ is reduced to $B(H_1) \oplus B(H_2)$. The operators in $B(H_1 \oplus H_2)$ that could distinguish between the states in Eqn. (14) are demoted to the status of “unobservable.” Thus, the algebra of observables is actually $B(H_1) \oplus B(H_2)$, and not $B(H_1 \oplus H_2)$.

Now, the algebraic approach provides two inversely related approaches to superselection rules.

1. First, we can follow the original “state space butchery” approach in a slightly more principled fashion: suppose that we are given some fields acting as operators on some big Hilbert space \mathcal{H} . Let \mathfrak{F} denote the algebra of field operators. [Here \mathfrak{F} is the analogue of the algebra $B(H_1 \oplus H_2)$, and \mathcal{H} is the analogue of $H_1 \oplus H_2$. In this case, however, we are not given an *a priori* decomposition of \mathcal{H} into a direct sum.] Suppose that we are also given a gauge group G that acts on the fields. We then *define* the observables as the gauge invariant fields. Let \mathfrak{A} denote the algebra of observables. We also define the physical pure states as those vectors in \mathcal{H} that transform irreducibly under the gauge group. A series of mathematical results (partially described in Section 9) shows then that \mathcal{H} decomposes into a direct sum $\bigoplus \mathcal{H}_\xi$ of subspaces that transform irreducibly under the gauge group; and each subspace \mathcal{H}_ξ corresponds to an irreducible representation of the algebra of observables \mathfrak{A} . We sketch this “top down” approach to superselection rules

³Superselection rules are also of foundational interest because they have been thought to help with the measurement problem — see e.g. [Beltrametti and Cassinelli, 1981, p. 74], [Landsman, 1995], [van Fraassen, 1991, pp. 264–272] — and more generally because of their connection with the emergence of a classical realm [Giulini, 2003]. However, we do not take up those specific issues in this Chapter.

in Section 9.

2. Instead of beginning with the field algebra \mathfrak{F} and deriving the superselection structure (i.e. the set of physically interesting representations of the algebra \mathfrak{A} of observables), we can begin with \mathfrak{A} and consider its set of physical representations. What is a ‘physical’ representation of \mathfrak{A} ? According to the criterion proposed by Doplicher, Haag, and Roberts (the DHR selection criterion), the physical representations are those that differ observably from the vacuum representation only in local regions. In this case, we still have the notion of superselection sectors, but we do not yet have a notion of fields or of a gauge group. It is not immediately clear that we have enough structure to explain the phenomena.

However, it is at this point that the deep mathematical analysis begins. First, one proves that the category of DHR representations corresponds precisely to the set Δ of localized transportable endomorphisms of the observable algebra \mathfrak{A} (see Section 8.2). Second, one proves that the set Δ naturally has the structure of a symmetric tensor $*$ -category (see Section 8). Finally, the Doplicher-Roberts Reconstruction Theorem shows that the unobservable fields \mathfrak{F} and gauge group G can be uniquely reconstructed from the category Δ .

The following sections outline some of the most important insights that have been gained in the study of superselection rules, and how this analysis bears on the foundational questions about the role of inequivalent representations. In short, our conclusion is that inequivalent representations are *not* irrelevant, and nor are they a problem. Rather, it is the *structure of the category of representations* that provides the really interesting theoretical content of QFT.

7.2 Minimal assumptions about the algebra of observables

For our discussion of superselection theory we need only a considerably pared down set of assumptions about the net of observable algebras. So, we now effectively cancel all assumptions we made about the net in Section 2. We begin with a *tabula rasa*, and add only those assumptions that we will need in the following sections.

By calling \mathfrak{A} a “net,” we are assuming that if $O_1 \subseteq O_2$ then $\mathfrak{A}(O_1) \subseteq \mathfrak{A}(O_2)$. But we do not promote this to the status of an Assumption.

ASSUMPTION 145 (Microcauality). If O_1 and O_2 are spacelike separated then $[\mathfrak{A}(O_1), \mathfrak{A}(O_2)] = \{0\}$.

ASSUMPTION 146 (Property B). The net $O \rightarrow \mathfrak{R}_0(O) \equiv \pi_0(\mathfrak{A}(O))''$ of von Neumann algebras satisfies property B, where (\mathcal{H}_0, π_0) is the GNS representation of \mathfrak{A} induced by ω_0 .

ASSUMPTION 147 (Duality). The pair (\mathfrak{A}, ω_0) satisfies *Haag duality*, i.e.

$$\pi_0(\mathfrak{A}(O'))' = \pi_0(\mathfrak{A}(O))'',$$

for each double cone O , where (\mathcal{H}_0, π_0) is the GNS representation of \mathfrak{A} induced by ω_0 .

ASSUMPTION 148 (Separability). The vacuum Hilbert space \mathcal{H}_0 is separable.

ASSUMPTION 149 (Nontriviality). For each double cone O , $\pi_0(\mathfrak{A}(O))$ contains an operator that is not a multiple of the identity; i.e. $\pi_0(\mathfrak{A}(O)) \neq \mathbb{C}I$.

A few remarks on these assumptions: (i) The first assumption is about the net \mathfrak{A} , but the remaining assumptions apply to a pair (\mathfrak{A}, ω_0) , where \mathfrak{A} is the quasilocal algebra and ω_0 is some fixed state. (ii) The duality assumption states that not only are the observables in $\mathfrak{R}_0(O')$ compatible with the observables in $\mathfrak{R}_0(O)$, but that $\mathfrak{R}_0(O')$ contains *all* observables that are compatible with the collection of observables in $\mathfrak{R}_0(O)$. We will assume in the following two sections (on DHR superselection theory) that the net \mathfrak{A} satisfies duality relative to some privileged vacuum state ω_0 . But, it does not follow from this that the net satisfies duality relative to every physical representation. In fact, a representation satisfies duality iff that sector has normal (Bose/Fermi) statistics; and every representation satisfies duality iff the gauge group is abelian. (iii) Duality in the vacuum sector is equivalent to the non-existence of spontaneously broken gauge symmetries. For the case of broken symmetries, we would impose a weaker requirement: essential duality. cf. Section 10.7. (iv) The separability assumption will only be invoked once — to show all superselection sectors induced by local fields are strongly locally equivalent (Prop. 243).

To be clear, note that we are now making no assumptions about the following: (i) No assumptions about the action of spacetime symmetries (e.g. translation symmetries, Lorentz symmetries) on the algebra \mathfrak{A} ; (ii) No assumptions to the effect that the vacuum state ω_0 is translation invariant; (iii) No assumptions about the action of spacetime symmetries on the vacuum Hilbert space; (iv) No assumptions about the spectrum condition.

8 THE CATEGORY Δ OF LOCALIZED TRANSPORTABLE ENDOMORPHISMS

In this Section we study the category $\Delta(\mathfrak{A})$ of localized transportable endomorphisms of the observable algebra \mathfrak{A} . Since the physical motivation for this study might not be initially clear, we pause to note the relation between endomorphisms and representations.

Suppose that π_0 is a fixed representation of \mathfrak{A} of some physical significance — e.g. the vacuum representation. Then for any endomorphism ρ of \mathfrak{A} , the composition $\pi_0 \circ \rho$ is also a representation of \mathfrak{A} . Thus, endomorphisms of \mathfrak{A} correspond naturally to representations of \mathfrak{A} , and we can hope to gain insight into the structure of the representations of \mathfrak{A} by studying the endomorphisms of \mathfrak{A} . However, the set $\text{End}\mathfrak{A}$ of endomorphisms of \mathfrak{A} has more intrinsic structure than the set $\text{Rep}\mathfrak{A}$ of representations of \mathfrak{A} — e.g., there is a product (viz. composition) operation on $\text{End}\mathfrak{A}$, and some endomorphisms have inverses. Thus, besides the traditional notions of equivalence and disjointness of representations, there are additional relations of physical importance on the set of representations of the form $\pi_0 \circ \rho$ with $\rho \in \text{End}\mathfrak{A}$.

If the problem of Lagrangian QFT is that there is only one Hilbert space, the problem of AQFT is that there are too many Hilbert spaces! Surely, not all of the representations of \mathfrak{A} are physical. In Section 9, we look at the problem from a more traditional point of view. In particular, we begin with a field algebra \mathfrak{F} of operators acting on a Hilbert space \mathcal{H} , and a gauge group G of unitary operators on \mathcal{H} . (We may suppose that G is the image of some representation of a fundamental symmetry group, e.g. $SU(2)$.) We also suppose that \mathcal{H} contains a vacuum state Ω . We then *define* the observable algebra \mathfrak{A} as the *gauge invariant fields*. But then we are again in the domain of AQFT: we have a reducible representation π of \mathfrak{A} on \mathcal{H} , and the irreducible subrepresentations of π are the *superselection sectors that can be reached from the vacuum sector by the action of local (unobservable) fields*. Not all representations of \mathfrak{A} appear in the decomposition of π — those that do not are surplus structure. However, all representations that appear in the decomposition of π are of the form $\pi_0 \circ \rho$, with ρ an endomorphism from the category $\Delta(\mathfrak{A})$! So, the motivation for studying these endomorphisms is that they correspond to representations that arise in this traditional, physically motivated way by acting on the vacuum representation with (unobservable) fields.⁴

There is yet another motivation for studying the DHR category: we want to understand the nature of gauge symmetries, and the DR Reconstruction Theorem provides crucial insight. In particular, the Theorem shows that DHR categories are in duality (in a mathematically precise sense) with compact groups. So, wherever there is a compact group, there is a DHR category, and vice versa. The study

⁴The DHR representations do *not* include those that can be reached from the vacuum by *nonlocal* fields, and so the domain of DHR superselection theory does not include theories with long range forces. But the case of local fields is already complicated enough, and is good training for the more general case.

of DHR categories and the study of compact gauge groups are one and the same; or, to take a more controversial stance, the structure of the category of physical representations of \mathfrak{A} explains why there is a compact gauge group (see [Roberts, 1975]).

We now define the category $\Delta = \Delta(\mathfrak{A})$ and uncover some of its natural structures. As stated above, the objects of our category Δ will be a subset of the set $\text{End}\mathfrak{A}$ of $*$ -endomorphisms of \mathfrak{A} .

DEFINITION 150. Let ρ be a $*$ -endomorphism of \mathfrak{A} , i.e. $\rho : \mathfrak{A} \rightarrow \mathfrak{A}$ is a $*$ -homomorphism (not necessarily surjective). Let O be a double cone in (Minkowski) spacetime. Then ρ is said to be *localized* in O just in case $\rho(A) = A$, for all $A \in \mathfrak{A}(O')$, where O' is the spacelike complement of O . We say that ρ is *localized* just in case there is a double cone O in which it is localized.

NOTE 151. By definition, a localized endomorphism satisfies $\rho(I) = I$, where I is the identity in \mathfrak{A} .

DEFINITION 152. If ρ is localized in O then ρ is said to be *transportable* just in case for any other double cone O_1 , there is a morphism ρ_1 localized in O_1 and a unitary operator $U \in \mathfrak{A}$ such that $U\rho(A) = \rho_1(A)U$ for all $A \in \mathfrak{A}$.

DEFINITION 153. For each double cone $O \in \mathcal{K}$, we let $\Delta(O)$ denote the set of transportable morphisms that are localized in O , and we let $\Delta = \bigcup_{O \in \mathcal{K}} \Delta(O)$. Elements of Δ are the *objects* of the DHR category.

We must now define arrows between the objects.

DEFINITION 154. Let $\rho, \rho' \in \Delta$. We define the set $\text{Hom}(\rho, \rho')$ of arrows between ρ and ρ' as follows:

$$\text{Hom}(\rho, \rho') := \{T \in \mathfrak{A} : T\rho(A) = \rho'(A)T, \forall A \in \mathfrak{A}\}.$$

If $T \in \text{Hom}(\rho, \rho')$ and $S \in \text{Hom}(\rho', \sigma)$ then we define $S \circ T = ST$, which is obviously in $\text{Hom}(\rho, \sigma)$.

Obviously, the identity $I \in \mathfrak{A}$ functions as the identity arrow for all objects; i.e. $I = \text{id}_\rho \in \text{End}(\rho)$ for all $\rho \in \text{Obj}(\Delta)$. Occasionally, we will write I_ρ to indicate that we are considering I as the identity of $\text{End}(\rho)$.

LEMMA 155. *Suppose that $\rho_i \in \Delta(O_i)$ for $i = 1, 2$, and that $T \in \text{Hom}(\rho_1, \rho_2)$. Then for any double cone O containing $O_1 \cup O_2$, we have $T \in \mathfrak{A}(O)$.*

Proof. Let $B \in \mathfrak{A}(O')$. Then

$$TB = T\rho_1(B) = \rho_2(B)T = BT.$$

Hence $T \in \mathfrak{A}(O')$. By duality in the vacuum sector, $T \in \mathfrak{A}(O)$. ■

PROPOSITION 156. *With the definition of hom-sets given above, Δ is a category.*

Proof. Completely straightforward. ■

So, we have shown that Δ is a category. In the remainder of this Section, we uncover more structure on Δ . We first show that Δ is a C^* -category; this involves showing that Δ has direct sums (an \oplus operation), subobjects, and the hom-sets of Δ are vector spaces with a $*$ -operation and norm $\|\cdot\|$ that obey the appropriate analogue of the C^* -algebra norm property. We then drop reference to the norms on the hom-sets, and show that there is a product operation \otimes on Δ such that (Δ, \otimes, ι) is a tensor $*$ -category.

DEFINITION 157. A category \mathcal{C} is said to be a *linear category* over the complex field \mathbb{C} , or a \mathbb{C} -linear category, just in case for all $X, Y \in \text{Obj}(\mathcal{C})$, $\text{Hom}(X, Y)$ is a complex vector space, and the composition \circ of morphisms is bilinear. When speaking of \mathbb{C} -linear categories, we stipulate that all functors should be \mathbb{C} -linear.

DEFINITION 158. A *$*$ -operation* on a \mathbb{C} -linear category \mathcal{C} is a map which assigns to an arrow $s \in \text{Hom}(X, Y)$ another arrow $s^* \in \text{Hom}(Y, X)$. This map has to be antilinear, involutive ($(s^*)^* = s$), and contravariant ($((s \circ t)^* = t^* \circ s^*)$). A $*$ -operation is *positive* iff $s^* \circ s = 0$ implies $s = 0$. A *$*$ -category* is a \mathbb{C} -linear category with a positive $*$ -operation.

NOTE 159. If \mathcal{C} is a $*$ -category, then for each $X \in \text{Obj}(\mathcal{C})$, $\text{End}(X)$ is a $*$ -algebra.

DEFINITION 160. A $*$ -category is called a *C^* -category* if for all $X, Y \in \text{Obj}(\mathcal{C})$, there is a norm $\|\cdot\|_{X,Y}$ on $\text{Hom}(X, Y)$ such that $(\text{Hom}(X, Y), \|\cdot\|_{X,Y})$ is a Banach space and

$$\begin{aligned} \|s \circ t\|_{X,Z} &\leq \|s\|_{Y,Z} \cdot \|t\|_{X,Y}, & \forall s \in \text{Hom}(Y, Z), \forall t \in \text{Hom}(X, Y) \\ \|s^* \circ s\|_{X,X} &= \|s\|_{X,Y}^2, & \forall s \in \text{Hom}(X, Y). \end{aligned}$$

We borrow some definitions from the theory of $*$ -algebras.

DEFINITION 161. Let \mathcal{C} be a $*$ -category. An arrow $f \in \text{Hom}(X, Y)$ is said to be an *isometry* just in case $f^* \circ f = \text{id}_X$. An arrow $f \in \text{Hom}(X, Y)$ is said to be *unitary* just in case f and f^* are isometries. An arrow $p \in \text{End}(Y) = \text{Hom}(Y, Y)$ is said to be a *projection* if $p = p^*$ and $p \circ p = p$.

NOTE 162. If $s \in \text{Hom}(Y, X)$ is an isometry then the arrow $p \equiv s \circ s^* \in \text{End}(X)$ is a projection.

DEFINITION 163. Let \mathcal{C} be a $*$ -category. If $X, Y \in \text{Obj}(\mathcal{C})$, then X is said to be a *subobject* of Y just in case there is an isometry $f \in \text{Hom}(X, Y)$. (Roughly speaking, there is an isometric embedding of X into Y .) The $*$ -category \mathcal{C} is said to *have subobjects* just in case for each $Y \in \text{Obj}(\mathcal{C})$ and projection $g \in \text{End}(Y)$, there is an $X \in \text{Obj}(\mathcal{C})$ and an isometry $f \in \text{Hom}(X, Y)$ such that $f \circ f^* = g$. The $*$ -category \mathcal{C} is said to *have direct sums* just in case for any two objects X, Y in \mathcal{C} , there is an object Z in \mathcal{C} and isometries $f \in \text{Hom}(X, Z)$, $g \in \text{Hom}(Y, Z)$ such that $f \circ f^* + g \circ g^* = \text{id}_Z$.

We begin by verifying that the DHR category Δ is a $*$ -category, i.e. the hom sets are vector spaces over \mathbb{C} , and there is a positive $*$ -operation.

LEMMA 164. *The DHR category Δ is a $*$ -category. That is, if $\rho, \sigma \in \text{Obj}(\Delta)$, then $\text{Hom}(\rho, \sigma)$ is a vector space over \mathbb{C} with the operations inherited from \mathfrak{A} (which is a vector space over \mathbb{C}), and the composition of arrows is bilinear. Furthermore, the $*$ -operation inherited from \mathfrak{A} is antilinear, involutive, contravariant, and positive.*

Proof. Completely straightforward. ■

PROPOSITION 165. *The DHR category Δ has direct sums.*

Proof. Let $\rho_1 \in \Delta(O_1)$, and let $\rho_2 \in \Delta(O_2)$. Choose a double cone O such that $(O_1 \cup O_2)^- \subseteq O$. Let E be a projection in $\mathfrak{A}(O_1)$. By property B, there are isometries $V_1, V_2 \in \mathfrak{A}(O)$ such that $V_1V_1^* + V_2V_2^* = I$. Define $\rho : \mathfrak{A} \rightarrow \mathfrak{A}$ by

$$\rho(A) = V_1\rho_1(A)V_1^* + V_2\rho_2(A)V_2^*, \quad \forall A \in \mathfrak{A}.$$

Since $V_iV_j = \delta_{ij}I$, and $\sum_i V_iV_i^* = I$, it follows that ρ is a morphism. Since ρ_1, ρ_2 are localized in O , and $V_1, V_2 \in \mathfrak{A}(O)$, it follows that ρ is localized in O .

To see that ρ is transportable, let \tilde{O} be another double cone. Since the ρ_i are transportable, there are endomorphisms ρ'_i localized in \tilde{O} , and unitary operators $U_i \in \text{Hom}(\rho_i, \rho'_i)$. As before, choose isometries V'_1, V'_2 in $\mathfrak{A}(\tilde{O})$, and set $\rho' = V'_1\rho'_1V'^*_1 + V'_2\rho'_2V'^*_2$. Then ρ' is localized in \tilde{O} and

$$V'_1U_1V_1^* \in \text{Hom}(\rho, \rho'), \quad V'_2U_2V_2^* \in \text{Hom}(\rho, \rho').$$

If we set $W = V'_1U_1V_1^* + V'_2U_2V_2^*$, then $W \in \text{Hom}(\rho, \rho')$ since it is a vector space. Furthermore,

$$\begin{aligned} W^*W &= [V'_1U_1V_1^* + V'_2U_2V_2^*]^*[V'_1U_1V_1^* + V'_2U_2V_2^*] \\ &= [V_1U_1^*V'^*_1 + V_2U_2^*V'^*_2][V'_1U_1V_1^* + V'_2U_2V_2^*] \\ &= V_1V_1^* + V_2V_2^* = I, \end{aligned}$$

and similarly for WW^* . Therefore W is a unitary operator in $\text{Hom}(\rho, \rho')$, showing that ρ is transportable. ■

DEFINITION 166. If $\rho_1, \rho_2 \in \Delta$, we denote their direct sum by $\rho_1 \oplus \rho_2$.

PROPOSITION 167. *The DHR category Δ has subobjects.*

Proof. Let $\rho \in \Delta(O)$, and let E be a projection in $\text{End}(\rho)$; i.e. $E\rho(A) = \rho(A)E$, for all $A \in \mathfrak{A}$. Then for all $A \in \mathfrak{A}(O')$,

$$EA = E\rho(A) = \rho(A)E = AE.$$

Therefore, by duality in the vacuum sector $E \in \mathfrak{A}(O)$. Choose O_1 such that $O^- \subseteq O_1$. By property B, there is an isometry $V \in \mathfrak{A}(O_1)$ such that $VV^* = E$. Now define $\rho' : \mathfrak{A} \rightarrow \mathfrak{A}$ by

$$\rho'(A) = V^*\rho(A)V, \quad \forall A \in \mathfrak{A}.$$

The isometry V embeds ρ' into ρ . Indeed,

$$\rho'(A)V^* = V^*\rho(A)VV^* = V^*\rho(A)E = V^*\rho(A).$$

and V is an isometry in $\text{Hom}(\rho', \rho)$ such that $VV^* = E \in \text{End}(\rho)$.

To see that ρ' is transportable, suppose that O_2 is an arbitrary double cone. Choose a double cone O_3 such that $O_3^- \subseteq O_2$. Since ρ is transportable, there is a morphism σ localized in O_3 and a unitary $U \in \text{Hom}(\rho, \sigma)$. It then follows that $U\text{End}(\rho)U^* = \text{End}(\sigma)$, thus $E' = UEU^*$ is a projection in $\text{End}(\sigma)$. Using property B, there is an isometry $V' \in \mathfrak{A}(O_1)$ such that $V'V'^* = E'$. Let $\sigma' = V'^*\sigma V'$. Clearly σ' is localized in O_1 , and $W = V'^*UV \in \text{Hom}(\rho', \sigma')$. Finally, W is unitary:

$$\begin{aligned} W^*W &= V^*U^*V'V'^*UV = V^*U^*E'UV \\ &= V^*EV = V^*VV^*V = I, \end{aligned}$$

and similarly for WW^* . Thus σ' is equivalent to ρ' . Since O_2 was an arbitrary double cone, $\rho' \in \Delta$. ■

DEFINITION 168. Suppose that \mathcal{C} is a \mathbb{C} -linear category. An object X in \mathcal{C} is said to be *irreducible* if it is nonzero and $\text{End}(X) = \mathbb{C}\text{id}_X$.

NOTE 169. Let ι be the identity endomorphism of \mathfrak{A} . Then $\iota \in \text{Obj}(\Delta)$, and since the vacuum representation of \mathfrak{A} is irreducible, ι is an irreducible object.

We now define a bifunctor $\otimes = (\otimes, \times)$ on the DHR category Δ , and verify that (Δ, \otimes, ι) is a tensor $*$ -category. But first we recall the pertinent definitions.

DEFINITION 170. A *bifunctor* on a category \mathcal{C} consists of two mappings $F : \text{Obj}(\mathcal{C}) \times \text{Obj}(\mathcal{C}) \rightarrow \text{Obj}(\mathcal{C})$ and $F : \text{Hom}(\mathcal{C}) \times \text{Hom}(\mathcal{C}) \rightarrow \text{Hom}(\mathcal{C})$, such that for $s \in \text{Hom}(X, Y)$ and $t \in \text{Hom}(X', Y')$, $F(s, t) \in \text{Hom}(F(X, X'), F(Y, Y'))$, and

$$\begin{aligned} F(s_1 \circ s_2, t) &= F(s_1, t) \circ F(s_2, t), \\ F(s, t_1 \circ t_2) &= F(s, t_1) \circ F(s, t_2), \\ F(\text{id}_X, \text{id}_{X'}) &= \text{id}_{F(X, X')}. \end{aligned}$$

If \mathcal{C} is a $*$ -category, then a bifunctor F is also required to be bilinear and to commute with the $*$ -operation. That is, for $s_i \in \text{Hom}(X, X')$, $t_i \in \text{Hom}(Y, Y')$ and $c \in \mathbb{C}$, we have

$$\begin{aligned} F(s_1 + s_2, t) &= F(s_1, t) + F(s_2, t), \\ F(s, t_1 + t_2) &= F(s, t_1) + F(s, t_2), \\ F(cs, t) &= cF(s, t) = F(s, ct), \end{aligned}$$


and

$$F(s, t)^* = F(s^*, t^*).$$

DEFINITION 171. Let $\otimes = (\otimes, \times)$ be a bifunctor on the category \mathcal{C} , and let $\mathbf{1} \in \text{Obj}(\mathcal{C})$. Then $(\mathcal{C}, \otimes, \mathbf{1})$ is said to be a *tensor category* just in case \otimes is associative up to a natural isomorphisms, and $\mathbf{1}$ is a two sided identity up to natural isomorphisms. The object $\mathbf{1}$ is called the *monoidal unit*. To be precise, to say that \otimes is ‘associative up to a natural isomorphisms’ means that for each $X, Y, Z \in \text{Obj}(\mathcal{C})$, there is an isomorphism $\alpha_{X,Y,Z} : X \otimes (Y \otimes Z) \rightarrow (X \otimes Y) \otimes Z$ that is ‘natural’ for all X, Y, Z ; i.e. if $s : X \rightarrow X'$ then

$$(15) \quad ((s \otimes \text{id}_Y) \otimes \text{id}_Z) \circ \alpha_{X,Y,Z} = \alpha_{X',Y,Z} \circ (s \otimes (\text{id}_Y \otimes \text{id}_Z)),$$

and similarly for Y and Z . Furthermore, α is required to make the pentagonal diagram commute:



To say that $\mathbf{1} \in \mathcal{C}$ is a two sided identity up to natural isomorphisms means that for each object $X \in \text{Obj}(\mathcal{C})$, there are isomorphisms $\lambda_X \in \text{Hom}(\mathbf{1} \otimes X, X)$ and $\rho_X \in \text{Hom}(X \otimes \mathbf{1}, X)$ such that:

1. λ_X and ρ_X are natural in X ; i.e. for any $s : X \rightarrow Y$,

$$(16) \quad s \circ \lambda_X = \lambda_Y \circ (\text{id}_{\mathbf{1}} \otimes s),$$

$$(17) \quad s \circ \rho_X = \rho_Y \circ (s \otimes \text{id}_{\mathbf{1}}).$$

In other words, the following two diagrams commute:

$$\begin{array}{ccc}
 \mathbf{1} \otimes X & \xrightarrow{\lambda_X} & X \\
 \downarrow \text{id}_{\mathbf{1}} \otimes s & & \downarrow s \\
 \mathbf{1} \otimes Y & \xrightarrow{\lambda_Y} & Y
 \end{array}
 \qquad
 \begin{array}{ccc}
 X \otimes \mathbf{1} & \xrightarrow{\rho_X} & X \\
 \downarrow s \otimes \text{id}_{\mathbf{1}} & & \downarrow s \\
 Y \otimes \mathbf{1} & \xrightarrow{\rho_Y} & Y
 \end{array}$$

2. λ_X and ρ_X make the triangular diagram commute:

$$\begin{array}{ccc}
 X \otimes (\mathbf{1} \otimes Y) & \xrightarrow{\alpha} & (X \otimes \mathbf{1}) \otimes Y \\
 \searrow \text{id}_X \otimes \lambda_Y & & \downarrow \rho_X \otimes \text{id}_Y \\
 & & X \otimes Y
 \end{array}$$

If \mathcal{C} is also a $*$ -category, there are two further requirements: (a.) the bifunctor \otimes must be compatible with the operations of $+$ and $*$ (as required in the definition of bifunctor), and (b.) the monoidal unit $\mathbf{1}$ must be irreducible, i.e. $\text{End}(\mathbf{1}) = \mathbb{C}\text{id}_{\mathbf{1}}$. For a C^* -category \mathcal{C} we require in addition that $\|s \times t\|_{X \otimes Y, X' \otimes Y'} \leq \|s\|_{X, X'} \cdot \|t\|_{Y, Y'}$.

Mac Lane’s coherence theorem shows that we can without danger ignore the natural isomorphisms α, λ , and ρ . That is, we can treat $X \otimes (Y \otimes Z)$ and $(X \otimes Y) \otimes Z$ as the *same* object, and we can treat $X, \mathbf{1} \otimes X$, and $X \otimes \mathbf{1}$ as the same object. To be more precise, we define:

DEFINITION 172. A tensor category \mathcal{C} is said to be *strict* if $\alpha_{X,Y,Z}, \lambda_X, \rho_X$ are identity morphisms for all $X, Y, Z \in \text{Obj}(\mathcal{C})$.

For example, the tensor category $(\text{Vect}, \otimes, \mathbb{C})$ of vector spaces is *not* strict, since e.g. $V \otimes \mathbb{C}$ is not literally the same vector space as V . On the other hand, a commutative monoid M can be thought of as a strict tensor category with one object and with arrows corresponding to elements of M . The coherence theorem can then be formulated as follows.

THEOREM 173 (Coherence Theorem). *Every tensor category is equivalent to a strict tensor category.*

Proof. See [Mac Lane, 1998]. ■

DEFINITION 174. If \mathcal{C} is a tensor category, then we let \mathcal{C}^{st} denote its strictification.

With these definitions in hand, we proceed now to define a bifunctor on Δ , and to verify that it satisfies all of the relevant properties. Our product \otimes of objects in Δ will be just the composition of endomorphisms.

PROPOSITION 175. *If $\rho, \sigma \in \text{Obj}(\Delta)$ then $\rho\sigma \in \text{Obj}(\Delta)$.*

Proof. It is clear that if ρ is localized in O_1 and σ is localized in O_2 , then $\rho\sigma$ is localized in any double cone that contains $O_1 \cup O_2$.

To see that $\rho\sigma$ is transportable, let O_3 be an arbitrary double cone. Since ρ and σ are transportable, there are $\rho', \sigma' \in \Delta(O_3)$ and unitaries $U \in \text{Hom}(\rho, \rho')$ and $V \in \text{Hom}(\sigma, \sigma')$. Then $\rho'\sigma'$ is localized in O_3 and $U\rho(V)$ is a unitary in $\text{Hom}(\rho\sigma, \rho'\sigma')$. Therefore, $\rho\sigma$ is transportable. ■

DEFINITION 176. Define $\otimes : \text{Obj}(\Delta) \times \text{Obj}(\Delta) \rightarrow \text{Obj}(\Delta)$ by $\rho \otimes \sigma = \rho\sigma$.

The product \times of arrows is slightly more complicated.

PROPOSITION 177. *If $S \in \text{Hom}(\rho, \rho')$ and $T \in \text{Hom}(\sigma, \sigma')$ then $S\rho(T) \in \text{Hom}(\rho \otimes \sigma, \rho' \otimes \sigma')$.*

Proof. Since $S\rho(T) = \rho'(T)S$, it follows that for any $A \in \mathfrak{A}$,

$$\begin{aligned} (S\rho(T))\rho\sigma(A) &= S\rho(T\sigma(A)) = \rho'(T\sigma(A))S = \rho'(\sigma'(A)T)S \\ &= \rho'\sigma'(A)(\rho'(T)S) = \rho'\sigma'(A)(S\rho(T)). \end{aligned}$$

Therefore $S\rho(T) \in \text{Hom}(\rho\sigma, \rho'\sigma')$. ■

DEFINITION 178. Define $\times : \text{Hom}(\Delta) \times \text{Hom}(\Delta) \rightarrow \text{Hom}(\Delta)$ by: for $S \in \text{Hom}(\rho, \rho')$ and $T \in \text{Hom}(\sigma, \sigma')$, we set $S \times T = S\rho(T) \in \text{Hom}(\rho \otimes \sigma, \rho' \otimes \sigma')$.

In the remainder of this section, we verify that (Δ, \otimes, ι) is a tensor $*$ -category.

\otimes is a bifunctor on Δ

PROPOSITION 179. *For $S_1, S_2, T_1, T_2 \in \text{Obj}(\Delta)$, if the source of T_i is the target of S_i (so that $T_i \circ S_i$ is defined) then*

$$(T_1 \times T_2) \circ (S_1 \times S_2) = (T_1 \circ S_1) \times (T_2 \circ S_2).$$

Proof. Straightforward calculation. ■

We must now check that \times is compatible with $*$

PROPOSITION 180. *For all $S, T \in \text{Hom}(\Delta)$,*

$$(S \times T)^* = S^* \times T^*.$$

Proof. Straightforward calculation. ■

ι is a monoidal unit

For each $\rho \in \text{Obj}(\Delta)$, $\iota \otimes \rho$ must be naturally isomorphic to $\rho \otimes \iota$ and to ρ (as expressed in the monoidal unit diagrams). But in the present case, $\iota \otimes \rho = \rho \otimes \iota = \rho$, so this natural isomorphism holds trivially.

Natural associativity of \otimes

Next, the product operation $\otimes = (\otimes, \times)$ must be associative up to natural isomorphisms, as expressed by the pentagonal diagram. But this is trivial in the present case, because associativity holds strictly; that is:

PROPOSITION 181. *For all $\rho_1, \rho_2, \rho_3 \in \text{Obj}(\Delta)$,*

$$\rho_1 \otimes (\rho_2 \otimes \rho_3) = (\rho_1 \otimes \rho_2) \otimes \rho_3,$$

and for all $T_1, T_2, T_3 \in \text{Hom}(\Delta)$,

$$T_1 \times (T_2 \times T_3) = (T_1 \times T_2) \times T_3.$$

Proof. The first claim follows trivially from the fact that composition of endomorphisms is associative. The second claim can be verified by a straightforward calculation. ■

LEMMA 182. *(Δ, \otimes, ι) is a C^* -tensor category with the norms inherited from \mathfrak{A} .*

Proof. We must verify that $\text{Hom}(\rho, \sigma)$ is closed in the norm on \mathfrak{A} . But this follows immediately from the fact that

$$\text{Hom}(\rho, \sigma) = \{T \in \mathfrak{A} : T\rho(A) = \sigma(A)T, \forall A \in \mathfrak{A}\}.$$

It's clear that $\|s \circ t\| \leq \|s\| \|t\|$. Furthermore,

$$\|S \times T\| = \|S\rho(T)\| \leq \|S\| \cdot \|\rho(T)\| \leq \|S\| \cdot \|T\|.$$

■

To this point we have shown that (i): Δ is a C^* -category, and (ii): Δ is a tensor $*$ -category. The following five Subsections are not linearly ordered. Subsection 8.1 shows how to define the canonical braiding $\varepsilon_{\rho_1, \rho_2}$ on (Δ, \otimes, ι) such that it is a 'braided' tensor $*$ -category. Then in Subsection 8.2 we make good our claims about the motivation for studying the category Δ : we prove that there is a functorial correspondence between Δ and the category of representations that satisfy the DHR selection criterion. We then pick up some technical information about tensor $*$ -categories that is essential for the physical interpretation of the corresponding representations. In Subsection 8.3 we see how to define a notion of the 'dimension' of an object in a tensor $*$ -category, and we define the notion of 'conjugate' objects. In Subsection 8.4 we take a detour to talk about the relation of spacetime symmetries to the DHR representations. Finally, in Subsection 8.5 we give the intrinsic statistical classification of objects of Δ_f that corresponds to the intuitive distinction between Bosons and Fermions, or Bose fields and Fermi fields.

8.1 Δ is a braided tensor $*$ -category

In this Subsection we define the canonical braiding on Δ ; this gives us a grasp on what happens if we change the order in products, say $\rho \otimes \sigma$ versus $\sigma \otimes \rho$. We will also see that there is a remarkable connection between spacetime dimension and the properties of this braiding: if the spacetime has three or more dimensions, the braiding is a symmetry. We first recall the pertinent definitions.

DEFINITION 183. If $(\mathcal{C}, \otimes, \mathbf{1})$ is a tensor category then a *braiding* on \mathcal{C} is a family of isomorphisms

$$\{c_{X,Y} \in \text{Hom}(X \otimes Y, Y \otimes X) : X, Y \in \text{Obj}(\mathcal{C})\},$$

satisfying the following two conditions:

1. $c_{X,Y}$ is natural in X and Y ; i.e. for any $f \in \text{Hom}(X, X')$ and $g \in \text{Hom}(Y, Y')$,

$$(18) \quad (g \times f) \circ c_{X,Y} = c_{X',Y'} \circ (f \times g).$$

2. $c_{X,Y}$ makes the following two hexagonal diagrams commute:

$$\begin{array}{ccc} (X \otimes Y) \otimes Z & \xrightarrow{c_{X \otimes Y, Z}} & Z \otimes (X \otimes Y) \\ \downarrow \alpha^{-1} & & \downarrow \alpha \\ X \otimes (Y \otimes Z) & & (Z \otimes X) \otimes Y \\ \downarrow \text{id}_X \otimes c_{Y,Z} & & \downarrow c_{Z,X} \otimes \text{id}_Y \\ X \otimes (Z \otimes Y) & \xrightarrow{\alpha} & (X \otimes Z) \otimes Y \end{array}$$

$$\begin{array}{ccc} X \otimes (Y \otimes Z) & \xrightarrow{c_{X, Y \otimes Z}} & (Y \otimes Z) \otimes X \\ \downarrow \alpha & & \downarrow \alpha^{-1} \\ (X \otimes Y) \otimes Z & & Y \otimes (Z \otimes X) \\ \downarrow c_{X,Y} \otimes \text{id}_Z & & \downarrow \text{id}_Y \otimes c_{Z,X} \\ (Y \otimes X) \otimes Z & \xrightarrow{\alpha^{-1}} & Y \otimes (X \otimes Z) \end{array}$$

That is, suppressing the associativity isomorphisms, $c_{X \otimes Y, Z}$ is expressed in terms of $c_{X, Y}$ and $c_{X, Z}$ as:

$$c_{X \otimes Y, Z} = (\text{id}_Y \otimes c_{Z, X})^{-1} \circ (\text{id}_X \otimes c_{Y, Z}),$$

and $c_{X, Y \otimes Z}$ is expressed in terms of $c_{X, Y}$ and $c_{Z, X}$ as:

$$c_{X, Y \otimes Z} = (\text{id}_Y \otimes c_{Z, X})^{-1} \circ (c_{X, Y} \otimes \text{id}_Z).$$

DEFINITION 184. A braiding $c_{X, Y}$ is called a *symmetry* if $(c_{X, Y})^{-1} = c_{Y, X}$ for all $X, Y \in \text{Obj}(\mathcal{C})$.

DEFINITION 185. A tensor category with a privileged braiding (symmetry) is called a *braided (symmetric) tensor category*.

In order to find our braiding on Δ , we will need the following technical lemma.

LEMMA 186. *If $\rho \in \Delta(O_1)$ and $\sigma \in \Delta(O_2)$ where O_1 and O_2 are spacelike separated, then $\rho\sigma = \sigma\rho$.*

Proof. Since the union of $\{\mathfrak{A}(O) : O_1 \cup O_2 \subseteq O\}$ is dense in \mathfrak{A} , it suffices to show that $\rho\sigma(A) = \sigma\rho(A)$ whenever $A \in \mathfrak{A}(O)$ with $O_1 \cup O_2 \subseteq O$. Choose O_3, O_4 that are spacelike to O and such that $O_1 \cup O_3$ is spacelike to $O_2 \cup O_4$. (This may always be done, even in two dimensional spacetime.) Since ρ, σ are transportable, there are ρ', σ' localized in O_3 and O_4 respectively and unitary operators $U_1 \in \text{Hom}(\rho, \rho')$ and $U_2 \in \text{Hom}(\sigma, \sigma')$. Then

$$\sigma(A) = U_2\sigma'(A)U_2^* = U_2^*AU_2.$$

Furthermore, $U_2 \in \mathfrak{A}(O'_1)$ by duality in the vacuum sector. Hence $\rho(U_2) = U_2$, and

$$\rho\sigma(A) = U_2U_1AU_1^*U_2^*.$$

Since $U_2U_1 = U_1U_2$, it follows that $\rho\sigma(A) = \sigma\rho(A)$. ■

We will not be able to define the braiding $\varepsilon_{\rho_1, \rho_2}$ in one shot. Rather, we first define arrows

$$\varepsilon_{\rho_1, \rho_2}(U_1, U_2) \in \text{Hom}(\rho_1 \otimes \rho_2, \rho_2 \otimes \rho_1),$$

that depend on the choice of “spectator morphisms” $\tilde{\rho}_1, \tilde{\rho}_2$, and unitary intertwiners $U_i \in \text{Hom}(\rho_i, \tilde{\rho}_i)$. We will then show that this definition is independent of the spectator morphisms and unitary intertwiners. (But, interestingly, when spacetime is two dimensional, the definition depends on the choice of a spatial orientation.)

DEFINITION 187. Suppose that $\rho_1 \in \Delta(O_1)$ and $\rho_2 \in \Delta(O_2)$. Let \tilde{O}_1 and \tilde{O}_2 be spacelike separated double cones. Since ρ_1 and ρ_2 are transportable, there are $\tilde{\rho}_i \in \Delta(\tilde{O}_i)$ and unitary operators $U_i \in \text{Hom}(\rho_i, \tilde{\rho}_i)$. Thus $U_1 \times U_2 \in \text{Hom}(\rho_1 \otimes \rho_2, \tilde{\rho}_1 \otimes \tilde{\rho}_2)$, and $U_2^* \times U_1^* \in \text{Hom}(\tilde{\rho}_2 \otimes \tilde{\rho}_1, \rho_2 \otimes \rho_1)$. Since \tilde{O}_1 is spacelike to \tilde{O}_2 ,

Lemma 186 entails that $\tilde{\rho}_1 \otimes \tilde{\rho}_2 = \tilde{\rho}_2 \otimes \tilde{\rho}_1$. Thus, we may define $\varepsilon_{\rho_1, \rho_2}(U_1, U_2) \in \text{Hom}(\rho_1 \otimes \rho_2, \rho_2 \otimes \rho_1)$ by

$$(19) \quad \varepsilon_{\rho_1, \rho_2}(U_1, U_2) := (U_2 \times U_1)^* \circ (U_1 \times U_2) = \rho_2(U_1^*)U_2^*U_1\rho_1(U_2).$$

NOTE 188. Since endomorphisms preserve unitarity, $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is unitary.

To show that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is independent of U_1 and U_2 , we need the following Lemma, which shows that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ does not change under certain “perturbations” of U_1 and U_2 .

LEMMA 189. *For $i = 1, 2$, let $\rho_i \in \Delta(O_i)$, let \tilde{O}_1 and \tilde{O}_2 be spacelike separated, let $\tilde{\rho}_i \in \Delta(\tilde{O}_i)$, and let $U_i \in \text{Hom}(\rho_i, \tilde{\rho}_i)$. Then $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ depends only on neighborhoods of U_1, U_2 in the following sense: if W_1, W_2 are unitaries such that $W_1 \in \mathfrak{A}(\tilde{O}'_2)$, $W_2 \in \mathfrak{A}(\tilde{O}'_1)$, and $W_1W_2 = W_2W_1$, then*

$$\varepsilon_{\rho_1, \rho_2}(W_1U_1, W_2U_2) = \varepsilon_{\rho_1, \rho_2}(U_1, U_2).$$

Proof. We must show that

$$(20) \quad (W_2U_2 \times W_1U_1)^* \circ (W_1U_1 \times W_2U_2) = (U_2^* \times U_1^*) \circ (U_1 \times U_2).$$

For any two unitary operators, $W_1, W_2 \in \mathfrak{A}$, we have

$$W_iU_i \times W_jU_j = W_iU_i\rho_i(W_j)\rho_i(U_j) = W_i\rho'_i(W_j)(U_i \times U_j).$$

Since $W_1 \in \mathfrak{A}(\tilde{O}'_2)$ and $\tilde{\rho}_2$ is localized in \tilde{O}_2 , $\tilde{\rho}_2(W_1) = W_1$; and similarly, $\tilde{\rho}_1(W_2) = W_2$. Hence, the left hand side of Eqn. 20 becomes

$$\begin{aligned} [(U_2 \times U_1)^*\tilde{\rho}_2(W_1^*)W_2^*][W_1\tilde{\rho}_1(W_2)(U_1 \times U_2)] &= (U_2 \times U_1)^*W_1^*W_2^*W_1W_2(U_1 \times U_2) \\ &= (U_2 \times U_1)^*(U_1 \times U_2), \end{aligned}$$

where we used $W_1W_2 = W_2W_1$ for the second equality. ■

LEMMA 190. *Let $\rho_i \in \Delta(O_i)$, $i = 1, 2$, and let $T \in \text{Hom}(\rho_1, \rho_2)$. Then $T \in \mathfrak{A}(O)$ for any double cone containing $O_1 \cup O_2$.*

Proof. Let O be a double cone containing $O_1 \cup O_2$, and let $A \in \mathfrak{A}(O')$. Then $\rho_1(A) = \rho_2(A) = A$, and so

$$TA = T\rho_1(A) = \rho_2(A)T = AT.$$

Hence $T \in \mathfrak{A}(O)'$, and by duality in the vacuum sector, $T \in \mathfrak{A}(O)$. ■

Now we can show that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ depends only on the localization regions of the spectator morphisms.

PROPOSITION 191. *$\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is definable in terms of ρ_1, ρ_2 , and the regions \tilde{O}_1, \tilde{O}_2 ; and does not change if the latter are replaced by double cones $\tilde{\tilde{O}}_1, \tilde{\tilde{O}}_2$ such that $\tilde{O}_1 \subseteq \tilde{\tilde{O}}_1$ and $\tilde{O}_2 \subseteq \tilde{\tilde{O}}_2$.*

Proof. (1.) We show first that for a given pair $(\tilde{O}_1, \tilde{O}_2)$ of spacelike separated double cones, the definition $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is independent of spectator morphisms $(\tilde{\rho}_1, \tilde{\rho}_2)$, and unitary intertwiners (U_1, U_2) . So, suppose that $\tilde{\rho}_i \in \Delta(\tilde{O}_i)$, and $U'_i \in \text{Hom}(\rho_i, \tilde{\rho}_i)$. Let $W_i = U'_i U_i^* \in \text{Hom}(\tilde{\rho}_i, \tilde{\rho}_i)$, so that $U'_i = W_i U_i$. Since W_i has left and right support in \tilde{O}_i , $W_i \in \mathfrak{A}(\tilde{O}_i) \subseteq \mathfrak{A}(\tilde{O}_j)'$. Thus, $W_1 W_2 = W_2 W_1$, and the hypotheses of Lemma 189 are satisfied. Therefore $\varepsilon_{\rho_1, \rho_2}(U_1, U_2) = \varepsilon_{\rho_1, \rho_2}(U'_1, U'_2)$.

(2.) Now let $\tilde{\tilde{O}}_1$ and $\tilde{\tilde{O}}_2$ be double cones such that $\tilde{\tilde{O}}_1 \perp \tilde{\tilde{O}}_2$, and $\tilde{O}_i \subseteq \tilde{\tilde{O}}_i$, for $i = 1, 2$. Choose $\tilde{\tilde{\rho}}_i \in \Delta(\tilde{\tilde{O}}_i)$, and unitaries $U'_i \in \text{Hom}(\rho_i, \tilde{\tilde{\rho}}_i)$. But we also have $\tilde{\rho}_i \in \Delta(\tilde{O}_i) \subseteq \Delta(\tilde{\tilde{O}}_i)$. And the first part of the proof shows that for fixed support regions $(\tilde{\tilde{O}}_1, \tilde{\tilde{O}}_2)$, the definition of $\varepsilon_{\rho_1, \rho_2}$ is independent of the choice of spectator morphisms and unitary intertwiners. Therefore $\varepsilon_{\rho_1, \rho_2}(U_1, U_2) = \varepsilon_{\rho_1, \rho_2}(U'_1, U'_2)$. ■

NOTE 192. We can always choose the spectator morphisms to be localized in strictly spacelike separated regions. Indeed, given \tilde{O}_1 and \tilde{O}_2 that are spacelike separated, choose $\tilde{\tilde{O}}_1$ such that $(\tilde{\tilde{O}}_1)^- \subseteq \tilde{O}_1$. But then the Lemma implies (by switching $\tilde{\tilde{O}}_1$ with \tilde{O}_1 , and setting $\tilde{\tilde{O}}_2 = \tilde{O}_2$) that we get the same definition from using $\tilde{\tilde{O}}_1$ or \tilde{O}_1 . More generally, since ρ_1 is transportable, the regions \tilde{O}_i can be chosen arbitrarily small.

NOTE 193. The previous note shows that a definition of $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is always equivalent to a definition using spectator morphisms localized in strictly spacelike separated regions. That is, there is a neighborhood N of zero such that $\tilde{O}_1 + x \subseteq \tilde{\tilde{O}}'_2$ for all $x \in N$. Again, since \tilde{O}_1 and $\tilde{O}_1 + x$ are contained in a double cone $\tilde{\tilde{O}}_1 \subseteq \tilde{\tilde{O}}'_2$, the previous Lemma (applied twice) entails that the pairs $(\tilde{O}_1, \tilde{O}_2)$ and $(\tilde{O}_1 + x, \tilde{O}_2)$ yield the same definition of $\varepsilon_{\rho_1, \rho_2}$.

By also shrinking \tilde{O}_2 if necessary, and repeating the above construction, we see that for any vector x , the pairs $(\tilde{O}_1 + x, \tilde{O}_2 + x)$ and $(\tilde{O}_1, \tilde{O}_2)$ yield the same definition of $\varepsilon_{\rho_1, \rho_2}$.

NOTE 194. In what follows, by “one dimensional spacetime”, we mean one dimensional space with zero dimensional time. In this case, a double cone is just an open interval, and “spacelike separated” means disjoint.

PROPOSITION 195. *For spacetime of dimension two or less, $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is definable in terms of ρ_1, ρ_2 and the spatial orientation of \tilde{O}_1 with respect to \tilde{O}_2 . That is, $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is independent of the choice of $(\tilde{O}_1, \tilde{O}_2)$, subject to the constraint of having the same spatial orientation.*

Proof. Let $\tilde{O}_i, \tilde{\tilde{O}}_i$ be given such that $\tilde{O}_1 \perp \tilde{O}_2$, $\tilde{\tilde{O}}_1 \perp \tilde{\tilde{O}}_2$, and \tilde{O}_1 is oriented with respect to \tilde{O}_2 as $\tilde{\tilde{O}}_1$ is with respect to $\tilde{\tilde{O}}_2$. Recall that translations of $(\tilde{O}_1, \tilde{O}_2)$ do not change $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$; nor does replacement of \tilde{O}_i with a double cone either containing it or contained in it, and spacelike to \tilde{O}_j . But $(\tilde{O}_1, \tilde{O}_2)$ can be replaced by $(\tilde{\tilde{O}}_1, \tilde{\tilde{O}}_2)$ in a series of such moves. ■

DEFINITION 196. For spacetime of two dimensions or less, fix a spatial orientation, and use $O_1 < O_2$ to indicate that O_1 is to the left of O_2 .

LEMMA 197. For spacetimes of dimension two or less, if the spatial orientation of \tilde{O}_1 with respect to \tilde{O}_2 is the opposite of the spatial orientation of \tilde{O}_1 with respect to \tilde{O}_2 , then

$$\varepsilon_{\rho_1, \rho_2}(U_1, U_2) = [\varepsilon_{\rho_2, \rho_1}(U'_2, U'_1)]^*.$$

Proof. For defining $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$, we can choose $\tilde{O}_1 = O_1$, $\tilde{O}_2 < O_1$, $\tilde{\rho}_1 = \rho_1$, and $U_1 = I_{\rho_1} = I \in \text{Hom}(\rho_1, \rho_1)$. In this case, the definition simplifies to

$$\varepsilon_{\rho_1, \rho_2}(I, U_2) = U_2^* \rho_1(U_2), \quad (\tilde{O}_2 < \tilde{O}_1).$$

Using the same spectator morphisms, we have

$$\varepsilon_{\rho_2, \rho_1}(U_2, I) = \rho_1(U_2^*) U_2, \quad (\tilde{O}_2 < \tilde{O}_1).$$

This latter expression uses the opposite spatial orientation. By the definability of $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ in terms of spatial orientation (Prop. 195), we see that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2) = [\varepsilon_{\rho_2, \rho_1}(U'_2, U'_1)]^*$ when the opposite spatial orientations are used for the two definitions. ■

DEFINITION 198 (The Canonical Braiding on Δ). For spacetimes of dimension two or less, we implement the convention that $\varepsilon_{\rho_1, \rho_2} = \varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ with $\tilde{O}_2 < \tilde{O}_1$. The previous Lemma shows that if we define $\bar{\varepsilon}_{\rho_1, \rho_2}$ with the opposite convention, then $\bar{\varepsilon}_{\rho_1, \rho_2} = (\varepsilon_{\rho_2, \rho_1})^*$. For spacetimes of dimension three or more, we define $\varepsilon_{\rho_1, \rho_2} = \varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ with \tilde{O}_1 and \tilde{O}_2 spacelike separated.

We now verify that $\varepsilon_{\rho_1, \rho_2}$ is a braiding on (Δ, \otimes, ι) .

PROPOSITION 199. $\varepsilon_{\rho, \sigma}$ is a braiding on the DHR category (Δ, \otimes, ι) .

Proof. (1) We first show that $\varepsilon_{\rho, \sigma}$ is natural in ρ and σ . For this it suffices to show that if $T \in \text{Hom}(\rho, \rho')$ then

$$(21) \quad (I_\sigma \times T) \circ \varepsilon_{\rho, \sigma} = \varepsilon_{\rho', \sigma} \circ (T \times I_\sigma),$$

$$(22) \quad \varepsilon_{\sigma, \rho} \circ (I_\sigma \times T^*) = (T^* \times I_\sigma) \circ \varepsilon_{\sigma, \rho'}.$$

Let O_1, O_2, O_3 be double cones such that $\rho \in \Delta(O_1)$, $\sigma \in \Delta(O_2)$, and $\rho' \in \Delta(O_3)$. Choose a double cone O_4 that is spacelike to O_i , $i = 1, 2, 3$; and if the spacetime dimension is less than three, choose O_4 to the left of all three. Choose $\sigma' \in \Delta(O_4)$, and $U \in \text{Hom}(\sigma, \sigma')$. Then $\varepsilon(\rho, \sigma) = U^* \rho(U)$ and $\varepsilon(\rho', \sigma) = U^* \rho'(U)$. Since $T \in \text{Hom}(\rho, \rho') \subseteq \mathfrak{A}(O'_4)$, it follows that $\sigma'(T) = T$. Thus,

$$\sigma(T) U^* \rho(U) = U^* \sigma'(T) \rho(U) = U^* T \rho(U) = U^* \rho'(U) T.$$

This establishes Eqn. 21. The second equation can be established by a similar calculation.

(2) Now we show that $\varepsilon_{\rho,\sigma}$ makes the hexagonal diagrams commute. Since Δ is strict monoidal, we can omit the associativity isomorphisms. That is, it suffices to show that

$$(23) \quad \varepsilon_{\rho \otimes \sigma, \tau} = (\varepsilon_{\rho, \tau} \times I_\sigma) \circ (I_\rho \times \varepsilon_{\sigma, \tau}),$$

$$(24) \quad \varepsilon_{\rho, \sigma \otimes \tau} = (I_\sigma \times \varepsilon_{\rho, \tau}) \circ (\varepsilon_{\rho, \sigma} \times I_\tau).$$

Choose $\tau' \in \Delta$ such τ' is supported in a region that is spacelike to the support regions of ρ, σ, τ ; for spacetimes of one or two dimensions, choose the support region of τ' to the left. Let $U \in \text{Hom}(\tau, \tau')$. Then $\varepsilon_{\rho, \tau} = U^* \rho(U)$, $\varepsilon_{\sigma, \tau} = U^* \sigma(U)$, and $\varepsilon_{\rho \otimes \sigma, \tau} = U^* \rho \sigma(U)$. Furthermore,

$$U^* \rho(U) \rho[U^* \sigma(U)] = U^* \rho[U U^* \sigma(U)] = U^* \rho \sigma(U),$$

establishing Eqn. 23. The second equation is proven analogously. ■

PROPOSITION 200. *For spacetimes of dimension two or less, $\varepsilon_{\rho_1, \rho_2}$ is the unique braiding on (Δ, \otimes, ι) such that $\varepsilon_{\rho_1, \rho_2} = I$ when $\rho_i \in \Delta(O_i)$ with $O_2 < O_1$. For spacetimes of dimension three or more, $\varepsilon_{\rho_1, \rho_2}$ is the unique braiding on (Δ, \otimes, ι) such that $\varepsilon_{\rho_1, \rho_2} = I$ when $\rho_i \in \Delta(O_i)$ with O_1 and O_2 spacelike separated.*

Proof. Choosing \tilde{O}_2 to the left of O_1 we can set $\varepsilon_{\rho_1, \rho_2} = \varepsilon_{\rho_1, \rho_2}(I, U_2) = U_2^* \rho_1(U_2)$ where $U_2 \in \text{Hom}(\rho_2, \rho'_2)$. Now let c_{σ_1, σ_2} be another braiding on (Δ, \otimes, ι) such that $c_{\sigma_1, \sigma_2} = I$ whenever σ_1 is localized in a region to the right of the localization region of σ_2 . Then since c_{ρ_1, ρ_2} is natural in ρ_1 and ρ_2 and $c_{\rho_1, \rho'_2} = I$,

$$c_{\rho_1, \rho_2} = (U_2^* \times I_{\rho_1}) \circ c_{\rho_1, \rho'_2} \circ (I_{\rho_1} \times U_2) = U_2^* \rho_1(U_2) = \varepsilon_{\rho_1, \rho_2}.$$

The proof for the higher dimensional case is structurally identical. ■

PROPOSITION 201. *For spacetime of dimension three or more, $\varepsilon_{\rho_1, \rho_2} = (\varepsilon_{\rho_2, \rho_1})^{-1}$, hence $\varepsilon_{\rho_1, \rho_2}$ is a symmetry on (Δ, \otimes, ι) .*

Proof. We first show that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is independent of the choice $(\tilde{O}_1, \tilde{O}_2)$ of supports for the spectator morphisms. (Compare the proof of Prop. 195.) Since the spacetime has at least three dimensions, there is a sequence of double cones $O_i, i = 1, \dots, n$, such that: $O_1 = \tilde{O}_2$, for each $i, O_i \cup O_{i+1}$ is contained in a double cone spacelike to \tilde{O}_1 , and O_n has the opposite spatial orientation to \tilde{O}_1 as did \tilde{O}_2 . Applying Prop. 191 repeatedly, we conclude that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2) = \varepsilon_{\rho_1, \rho_2}(U_1, U'_2)$, where $\tilde{\rho}_2 \in \Delta(O_n)$ and $U'_2 \in \text{Hom}(\rho_2, \tilde{\rho}_2)$. Thus, $\varepsilon_{\rho_1, \rho_2}$ does not depend on the relative spatial orientation of \tilde{O}_1 and \tilde{O}_2 . Prop. 195 shows that $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ can depend on $(\tilde{O}_1, \tilde{O}_2)$ only through their relative spatial orientation. Therefore, $\varepsilon_{\rho_1, \rho_2}(U_1, U_2)$ is independent of $(\tilde{O}_1, \tilde{O}_2)$.

We can choose $\tilde{O}_1 = O_1$, $\tilde{O}_2 \perp O_1$, $\tilde{\rho}_1 = \rho_1$, and $U_1 = I_{\rho_1} = I \in \text{Hom}(\rho_1, \rho_1)$ so that

$$\varepsilon_{\rho_1, \rho_2} = U^* \rho_1(U_2).$$

But given the independence of $\varepsilon_{\rho_1, \rho_2}$ from the orientation of $(\tilde{O}_1, \tilde{O}_2)$, we also have

$$\varepsilon_{\rho_2, \rho_1} = \rho_1(U_2)^* U_2 = (\varepsilon_{\rho_1, \rho_2})^*.$$

Since $\varepsilon_{\rho_1, \rho_2}$ is unitary, $\varepsilon_{\rho_1, \rho_2} = (\varepsilon_{\rho_2, \rho_1})^{-1}$. ■

NOTE 202. The preceding Proposition is the first place where we invoked the dimension of the underlying spacetime. We will be clear when subsequent results depend on assumptions about dimension.

DEFINITION 203. Let $\varepsilon_\rho := \varepsilon(\rho, \rho) \in \text{End}(\rho \otimes \rho)$.

8.2 Relation between localized endomorphisms and representations

While the categories Δ and Δ_f defined in this section have very remarkable properties, their physical and philosophical relevance is certainly not obvious. We therefore relate the category Δ to a certain category of representations of the net \mathfrak{A} :

DEFINITION 204. Let $O \mapsto \mathfrak{A}(O)$ be a net of observables and $\pi_0 : \mathfrak{A} \rightarrow B(H_0)$ a vacuum representation. Then a DHR-representation (w.r.t. the vacuum representation π_0) is a $*$ -representation $\pi : \mathfrak{A} \rightarrow B(H)$ such that $\pi|_{\mathfrak{A}(O')} \cong \pi_0|_{\mathfrak{A}(O')}$ for any double cone O . I.e., upon restriction to $\mathfrak{A}(O')$, the representations π and π_0 are unitarily equivalent. The category whose objects are DHR-representations of \mathfrak{A} with bounded intertwining operators is denoted by $DHR(\mathfrak{A})$. It clearly is a C^* -category.

DEFINITION 205. Let \mathfrak{A} be a net that is Poincaré covariant w.r.t. the positive energy representation $U_0 : \mathcal{P} \rightarrow \mathcal{U}(H_0)$. A representation (H, π) of \mathfrak{A} is called covariant (with positive energy) if it is equipped with a strongly continuous unitary representation $U_\pi : \hat{\mathcal{P}} \rightarrow \mathcal{U}(H)$ (with $\text{spec} P^\mu \subset V_+$) of the universal covering of the Poincaré group such that $\text{Ad } U_\pi(h) \circ \pi = \pi \circ \alpha_h$ for all $h \in \hat{\mathcal{P}}$, where, omitting the covering map $\hat{\mathcal{P}} \rightarrow \mathcal{P}$ from the notation, $\alpha_h = \text{Ad } U_0(h)$.

Note that the definition implies that the representation space H of a DHR representation must have the same dimension as the vacuum Hilbert space H_0 .

PROPOSITION 206. *There is a functor $F : \Delta \rightarrow DHR(\mathfrak{A})$ of C^* -categories such that $F(\rho) = \pi_0 \circ \rho$ for the objects and $F(s) = \pi_0(s)$ for $s \in \text{Hom}_\Delta(\rho, \sigma)$ for morphisms. This functor is an equivalence.*

Proof. We first note that these definitions make sense: $\rho \in \text{Obj}(\Delta)$ maps \mathfrak{A} into itself and can therefore be composed with the representation π_0 , defining a new

representation. Furthermore, if S is an arrow in Δ , then Lemma 155 gives $S \in \mathfrak{A}$, thus $F(S) = \pi_0(S)$ makes sense. With $S \in \text{Hom}_\Delta(\rho, \rho')$ we have

$$\begin{aligned} F(S)F(\rho)(A) &= \pi_0(S)\pi_0(\rho(A)) = \pi_0(S\rho(A)) = \\ \pi_0(\rho'(A)S) &= F(\rho')(A)F(S) \quad \forall A \in \mathfrak{A}, \end{aligned}$$

thus $F(S) \in \text{Hom}(F(\rho), F(\rho'))$. Since id_ρ is the unit of \mathfrak{A} we have $F(\text{id}_\rho) = I_{H_0} = \text{id}_{F(\rho)}$. The property $F(S \circ T) = F(S) \circ F(T)$ is obvious. Since π_0 is faithful, F is faithful. We must show that the representation $F(\rho) = \pi_0 \circ \rho$ satisfies the DHR criterion. Since $\rho \in \Delta$ is transportable, for every double cone O there exist $\rho_O \in \Delta$ localized in O and a unitary $U_O : \rho \rightarrow \rho_O$. Since ρ_O is localized in O , the representation $F(\rho_O) = \pi_0 \circ \rho_O$ coincides with π_0 on $\mathfrak{A}(O')$. Since $F(U_O) : F(\rho) \rightarrow F(\rho_O)$ is unitary, we have

$$F(\rho)|_{\mathfrak{A}(O')} \cong F(\rho')|_{\mathfrak{A}(O')} = \pi_0|_{\mathfrak{A}(O')},$$

implying $F(\rho) = \pi_0 \circ \rho \in \text{DHR}(\mathfrak{A})$. Now let $\rho, \rho' \in \text{Obj}(\Delta)$ and $\tilde{S} \in \text{Hom}(F(\rho), F(\rho'))$. If O is a double cone containing the localization regions of ρ, ρ' ,

$$\tilde{S}\pi_0(A) = \tilde{S}\pi_0(\rho(A)) = \tilde{S}F(\rho)(A) = F(\rho')(A)S = \pi_0(\rho'(A))\tilde{S} = \pi_0(A)\tilde{S}$$

for all $A \in \mathfrak{A}(O')$. Therefore, by Haag duality for π_0 , $\tilde{S} \in \pi_0(\mathfrak{A}(O'))' = \pi_0(\mathfrak{A}(O))$. Thus there exists $S \in \text{Hom}_\Delta(\rho, \rho')$ such that $\tilde{S} = F(S)$. This proves that the functor F is full. Finally, let $\pi \in \text{DHR}(\mathfrak{A})$ be a DHR representation on a Hilbert space H . Choose any double cone O . Then the DHR criterion implies the existence of a unitary $U : H \rightarrow H_0$ such that $U\pi(A) = \pi_0(A)U$ for all $A \in \mathfrak{A}(O')$. Define a new representation π' of \mathfrak{A} on the vacuum Hilbert space H_0 by $\pi'(\cdot) = U\pi(\cdot)U^*$. By the very definition, we have $\pi'(A) = \pi_0(A)$ for all $A \in \mathfrak{A}(O')$. If now \hat{O} is any double cone containing O , and $A \in \mathfrak{A}(\hat{O}')$ and $B \in \mathfrak{A}(\hat{O})$ then

$$\pi'(B)\pi_0(A) = \pi'(BA) = \pi'(AB) = \pi'(A)\pi'(B) = \pi_0(A)\pi'(B),$$

implying $\pi'(\mathfrak{A}(\hat{O})) \subset \pi_0(\mathfrak{A}(\hat{O}'))' = \pi_0(\mathfrak{A}(O))$ by Haag duality for π_0 . Thus π' maps the quasilocal algebra \mathfrak{A} into $\pi_0(\mathfrak{A})$. Since π_0 is injective, we can define an endomorphism ρ of \mathfrak{A} by $\rho = \pi_0^{-1} \circ \pi'$. By construction, ρ is localized in O , and we have $\pi' = \pi_0 \circ \rho = F(\rho)$. Repeating the argument with a different double cone \tilde{O} , we see that ρ is transportable, thus $\rho \in \Delta$. Since $\pi \cong \pi' = F(\rho)$, we have proven that every DHR representation is unitarily equivalent to one of the form $F(\rho)$ where $\rho \in \Delta$. Thus the functor F is essentially surjective, and therefore, cf. Appendix A an equivalence of categories. ■

NOTE 207. The significance of Proposition 206 is twofold. On the one hand, it provides an interpretation of the category Δ in terms of a class of representations of \mathfrak{A} . If one accepts for a moment that the category $\text{DHR}(\mathfrak{A})$ is worth studying, the above equivalence is a powerful tool. Namely, it permits us to pull the symmetric monoidal structure on Δ over to $\text{DHR}(\mathfrak{A})$ – which as defined initially is just a category – such as to make the functor $F : \Delta \rightarrow \text{DHR}(\mathfrak{A})$ an equivalence of

symmetric tensor C^* -categories. But once this is understood, it is much more convenient to work just with the category Δ rather than with $DHR(\mathfrak{A})$, since the tensor structure on $DHR(\mathfrak{A})$ will not be strict.

As to the physical motivation of the DHR condition, we give three arguments:

1. By an increasing sequence of double cones we mean a sequence $O_1 \subset O_2 \subset \dots$ of double cones such that $\cup_i O_i = \mathbb{R}^d$ (typically $d = 4$). In the appendix of [Doplicher *et al.*, 1971], the following result (the converse of which is trivial) is proven:

THEOREM 208. *Let ω be a state on \mathfrak{A} such that*

$$\lim_{n \rightarrow \infty} \|(\omega - \omega_0)|\mathfrak{A}(O'_n)\| = 0$$

and the GNS-representation π_ω associated with ω satisfies property B. Then there is a double cone O such that $\pi_\omega|\mathfrak{A}(O) \cong \pi_0|\mathfrak{A}(O)$.

2. In Section 9 we will show that the DHR criterion is satisfied by superselection sectors that are connected to the vacuum sector by a field net satisfying Bose-Fermi commutation relations. (See Section 9 for precise definitions and statements.) Conversely, in Section 10 we will prove that every DHR representation with finite dimension arises in this way. Together these results imply that DHR superselection sectors are precisely those induced by (graded) local fields. We refer to Sections 9-10 for further discussion.
3. Let (H, π) be a Poincaré covariant representation (in the sense of Definition 205) of \mathfrak{A} such that H is separable and the spectrum $\text{sp } P^\mu \subset \mathbb{R}^d$ of the momentum operator P^μ has an isolated mass shell $\{p \mid p^2 = m^2\}$ at its bottom, where $m > 0$. (Such a representation is called a massive one-particle representation.) Then, as proven in [Buchholz and Fredenhagen, 1982], for every ‘spacelike cone’ \mathcal{C} one has a unitary equivalence $\pi|\mathfrak{A}(\mathcal{C}') \cong \pi_0|\mathfrak{A}(\mathcal{C}')$. (For the definition of spacelike cones cf. [Buchholz and Fredenhagen, 1982].) Despite the fact that this localization property is weaker than the one imposed by the DHR criterion, the category of representations that are localized in spacelike cones still can be equipped with a braided monoidal structure, cf. [Buchholz and Fredenhagen, 1982]. (The purely representation theoretic part of this theory was considerably simplified in [Doplicher and Roberts, 1990, Section 4].) In this theory, the dimension of spacetime must be $\geq 3 + 1$ in order for the braiding to be a symmetry! On the technical side the mathematical treatment is more complicated for the following reason: If π is a representation such that $\pi|\mathfrak{A}(\mathcal{C}') = \pi_0|\mathfrak{A}(\mathcal{C}')$, then Haag duality implies $\pi(\mathfrak{A}(\mathcal{C})) \subset \pi(\mathfrak{A}(\mathcal{C}))'$, but due to the weak closure the right hand side is not contained in the algebra \mathfrak{A} . The construction of a field net that we discuss in Section 10 can nevertheless be generalized to charges localized in spacelike cones, cf. [Doplicher and Roberts, 1990, Section 5]. On the grounds of the

cited results it seems evident that the cone-localized superselection sectors are physically better motivated than the more restrictive DHR sectors. The D(H)R theory expounded in Sections 7-10 remains useful as a technically easier ‘mathematical laboratory’.

8.3 Dimension theory in tensor *-categories

For any tensor *-category, we can define a notion of “conjugates.” The following is a simplified version of this definition for the case of a *strict* tensor *-category.

DEFINITION 209. Let \mathcal{C} be a strict tensor *-category and $X \in \text{Obj}(\mathcal{C})$. A *solution of the conjugate equations* is a triple $(\overline{X}, r, \bar{r})$, where $\overline{X} \in \text{Obj}(\mathcal{C})$ and $r : \mathbf{1} \rightarrow \overline{X} \otimes X$, $\bar{r} : X \otimes \overline{X} \rightarrow \mathbf{1}$ satisfy

$$\begin{aligned} (\bar{r}^* \otimes \text{id}_X) \circ (\text{id}_X \otimes r) &= \text{id}_X, \\ (r^* \otimes \text{id}_{\overline{X}}) \circ (\text{id}_{\overline{X}} \otimes \bar{r}) &= \text{id}_{\overline{X}}. \end{aligned}$$

A strict tensor *-category \mathcal{C} has *conjugates* if there is a solution of the conjugate equations for every $X \in \mathcal{C}$.

EXAMPLE 210. The definition of conjugates is exemplified in the (strictification of the) category $\text{Rep}_f G$ of finite dimensional representations of a compact group. In particular, it is well known that for each representation (H, π) of G , there is a conjugate representation $(\overline{H}, \bar{\pi})$ of G . (There are several different constructions of this conjugate representation; see e.g. [Simon, 1996, p. 30].) In terms of universal properties, $(\overline{H}, \bar{\pi})$ is the unique irreducible representation of G such that $(H \otimes \overline{H}, \pi \otimes \bar{\pi})$ contains a copy of the trivial representation of G .

NOTE 211. Suppose that $(\overline{X}, r, \bar{r})$ is a conjugate for X , and that the tensor unit $\mathbf{1}$ is irreducible. Then $r^* \circ r \in \text{End}(\mathbf{1}) = \mathbb{C}\text{id}_{\mathbf{1}}$. Thus up to a scalar, r is an isometry, and therefore $\mathbf{1}$ is a direct summand of $\overline{X} \otimes X$. Furthermore, as can be shown using the conjugate equations, the map $\text{End}(X) \rightarrow \text{Hom}(\mathbf{1}, \overline{X} \otimes X)$, defined by $s \mapsto (\text{id}_{\overline{X}} \otimes s) \circ r$, is an isomorphism of vector spaces. Therefore, if X is irreducible, the direct summand $\mathbf{1}$ appears with multiplicity 1 in $\overline{X} \otimes X$.

DEFINITION 212. Let \mathcal{C} be a tensor *-category and $X \in \text{Obj}(\mathcal{C})$. A solution $(\overline{X}, r, \bar{r})$ of the conjugate equations relative to X is called *normalized* if

$$r^* \circ r = \bar{r}^* \circ \bar{r},$$

and *standard* if

$$r^* \circ (\text{id}_{\overline{X}} \otimes a) \circ r = \bar{r}^* \circ (a \otimes \text{id}_{\overline{X}}) \circ \bar{r},$$

for all $a \in \text{End}(X)$.

NOTE 213. If X, Y have (standard) conjugates, then $X \otimes Y$ and $X \oplus Y$ also have (standard) conjugates. If an object has a conjugate, then it has a standard conjugate. For more details, see the appendix.

DEFINITION 214. If an object $X \in \text{Obj}(\mathcal{C})$ has a standard conjugate (\bar{X}, r, \bar{r}) , we define its dimension $d(X)$ by

$$d(X)\text{id}_{\mathbf{1}} = r^* \circ r.$$

If an object X does not have a conjugate, we formally say $d(X) = +\infty$.

NOTE 215. For all $X \in \text{Obj}(\mathcal{C})$, $d(X) \geq 0$. Furthermore, if $X, Y \in \text{Obj}(\mathcal{C})$ have conjugates then

$$d(\bar{X}) = d(X), \quad d(X \otimes Y) = d(X) \cdot d(Y), \quad d(X \oplus Y) = d(X) + d(Y),$$

and $d(\mathbf{1}) = 1$. (See the appendix for the discussion of these facts.)

DEFINITION 216. Let Δ be the DHR category. We define the full subcategory Δ_f of objects with finite dimension:

$$\text{Obj}(\Delta_f) = \{\rho \in \text{Obj}(\Delta) : d(\rho) < +\infty\}.$$

NOTE 217. By definition, Δ_f is a category with conjugates. It is closed under tensor products, direct sums and subobjects. In any C^* -tensor category with conjugates, the dimension of any object takes values in $[1, \infty)$, and in the interval $[1, 2]$ only values of the form $2 \cos(\pi/n)$, $n \geq 3$ can appear, cf. [Longo and Roberts, 1997]. In a symmetric C^* -tensor category, all dimensions are integers, as is proven in the Appendix.

PROPOSITION 218. *For each $X, Y \in \text{Obj}(\Delta_f)$, $\text{Hom}(X, Y)$ is a finite dimensional vector space. Every object $X \in \text{Obj}(\Delta_f)$ is a finite direct sum of irreducible objects; i.e. the category Δ_f is semisimple.*

Proof. See the appendix. ■

NOTE 219. There is an important connection, discovered by Longo [1989] and explored further in [Longo and Roberts, 1997], between the dimension of a DHR sector $\rho \in \Delta$ and subfactor theory. Among many other things, the latter associates to any inclusion $N \subset M$ of factors an index $[M : N] \in [1, \infty]$. In order to apply this theory to AQFT we need to assume (or prove) that the local von Neumann algebras $\mathfrak{A}(O)$ are factors. (This is automatic, e.g., in conformally covariant theories.) If $\rho \in \Delta$ is localized in O , it restricts to a normal $*$ -homomorphism of $\mathfrak{A}(O)$ into itself, giving rise to an inclusion $\rho(\mathfrak{A}(O)) \subset \mathfrak{A}(O)$. The index of this subfactor is related to the categorically defined dimension $d(\rho)$ by

$$(25) \quad [\mathfrak{A}(O) : \rho(\mathfrak{A}(O))] = d(\rho)^2.$$

Longo’s result allows to give a very direct formula for the dimension of (the localized endomorphisms associated to) a DHR representation. Namely, all endomorphisms $\rho \in \Delta$ for which $\pi \cong \pi_0 \circ \rho$ have the same categorical dimension, justifying to write $d(\pi)$, and for any double cone O we have

$$d(\pi) = [\pi(\mathfrak{A}(O'))' : \pi(\mathfrak{A}(O))]^{1/2}.$$

This is seen as follows: π is unitarily equivalent to a representation $\pi' = \pi_0 \circ \rho$, where $\rho \in \Delta$ is localized in O . Then the inclusion $\pi(\mathfrak{A}(O)) \subset \pi(\mathfrak{A}(O'))'$ is unitarily equivalent to $\pi'(\mathfrak{A}(O)) \subset \pi(\mathfrak{A}(O'))'$, which equals $\pi_0(\rho(\mathfrak{A}(O))) \subset \pi_0(\mathfrak{A}(O))$. Now the claim follows by Eqn. (25) and the fact that the index is invariant under unitary transformations: $[UMU^* : UNU^*] = [M : N]$.

Another comment seems in order: The categorical definition of dimension of an object requires the existence of a conjugate object. On the other hand, assuming factoriality of the local algebras, the expressions $[\mathfrak{A}(O) : \rho(\mathfrak{A}(O))]$ (for an endomorphism localized in O) and $[\pi(\mathfrak{A}(O'))' : \pi(\mathfrak{A}(O))]$ (whose independence of O follows from mild additional axioms) do not presuppose the existence of a conjugate. In fact, one can show that finiteness of these subfactor indices implies the existence of a conjugate DHR representation, cf. [Guido and Longo, 1992].

8.4 Covariant representations

Since we decided to work with the category Δ of localized transportable endomorphisms rather than directly with DHR representations, we need the following

DEFINITION 220. Let \mathfrak{A} be a Poincaré covariant net with covariant vacuum representation (H_0, π_0) . An endomorphism $\rho \in \Delta(\mathfrak{A})$ is called covariant if there exists a strongly continuous positive energy representation $\pi_\rho : \widehat{\mathcal{P}} \rightarrow U(H_0)$ such that

$$(26) \quad \text{Ad}U_\rho(h) \circ \pi_0 \circ \rho = \pi_0 \circ \rho \circ \beta_h \quad \forall h \in \widehat{\mathcal{P}}.$$

The full subcategory of $\Delta(\mathfrak{A})$ consisting of the covariant morphisms is denoted by $\Delta_c(\mathfrak{A})$.

NOTE 221. For $\rho \in \Delta$, $h \in \mathcal{P}$ we define $\rho_h = \beta_h \circ \rho \circ \beta_h^{-1}$. If ρ is localized in the double cone O then ρ_h is localized in hO . If $\rho \in \Delta_c$ then Eqn. (26) can be restated as

$$\text{Ad}(U(h)U_\rho(h)^*) \circ \pi_0 \circ \rho = \pi_0 \circ \beta_h \circ \rho \circ \beta_h^{-1} = \pi_0 \circ \rho_h \quad \forall h \in \widehat{\mathcal{P}}.$$

Since ρ and ρ_h are both localized, it follows that $X_\rho(h) \equiv U(h)U_\rho(h)^* \in \text{Hom}(\rho, \rho_h)$, thus $X_\rho(h) \in \mathfrak{A}$. This \mathfrak{A} -valued cocycle is very convenient since expressions like $\rho(U(h))$ don't make sense, whereas $\rho(X_\sigma(h))$ does. It satisfies the following cocycle equation:

$$\begin{aligned} X_\rho(gh) &= U(gh)U_\rho(gh)^* = U(g)U(h)U_\rho(h)^*U_\rho(g)^* \\ &= \beta_g(U(h)U_\rho(h)^*)U(g)U_\rho(g)^* = \beta_g(X_\rho(h))X_\rho(g). \end{aligned}$$

The same computation implies that, if $\rho \in \Delta$ and $h \mapsto X_\rho(h) \in \mathfrak{A}$ satisfies $X_\rho(gh) = \beta_g(X_\rho(h))X_\rho(g)$ for all $g, h \in \mathcal{P}$, then $U_\rho(h) := X_\rho(h)^*U(h)$ is a representation of \mathcal{P} and Eqn. 26 holds, i.e. $\rho \in \Delta_c$.

PROPOSITION 222. Δ_c is closed under tensor products, direct sums and subobjects.

Proof. Let $\rho, \rho' \in \Delta_c$ with associated cocycles $X_\rho, X_{\rho'}$. Then

$$(27) \quad X_{\rho\rho'}(h) = X_\rho(h) \otimes X_{\rho'}(h) = X_\rho(h)\rho(X_{\rho'}(h)) \in \text{Hom}(\rho \otimes \rho', \rho_h \otimes \rho'_h)$$

clearly satisfies the cocycle equation, thus $\rho\rho'$ is covariant. The proof for direct sums and subobjects is omitted, cf. [Roberts, 1990]. ■

If $T \in \text{Hom}(\rho, \rho')$ then

$$\beta_h(T)\rho_h(A) = \beta_h(T\rho\beta_h^{-1}(A)) = \beta_h(\rho'\beta_h^{-1}(A)T) = \rho'_h(A)\beta_h(T),$$

thus $\beta_h(T) \in \text{Hom}(\rho_h, \rho'_h)$.

Now we explore some consequences of finite dimensionality:

PROPOSITION 223. *Let $\rho, \rho' \in \Delta_{fc} := \Delta_f \cap \Delta_c$. Then*

- (i) *If $T \in \text{Hom}(\rho, \rho')$ then $TU_\rho(h) = U_{\rho'}(h)T$ for all $h \in \widehat{\mathcal{P}}$.*
- (ii) *Every $\rho \in \Delta_{fc}$ is covariant w.r.t. a unique representation U_ρ .*
- (iii) *If $\rho, \rho' \in \Delta_{fc}$ and $T \in \text{Hom}(\rho, \rho')$ then*

$$(28) \quad \beta_h(T)X_\rho(h) = X_{\rho'}(h)T \quad \forall h \in \widehat{\mathcal{P}}.$$

- (iv) *Δ_{fc} is closed under conjugates.*

Proof. (i) For $h \in \widehat{\mathcal{P}}$ define $T_h = U_{\rho'}(h)TU_\rho(h)^*$. For any $A \in \mathfrak{A}$ we have

$$\begin{aligned} T_h\rho(A) &= U_{\rho'}(h)TU_\rho(h)^*\rho(A) = U_{\rho'}(h)T\rho(\beta_h^{-1}(A))U_\rho(h)^* \\ &= U_{\rho'}(h)\rho'(\beta_h^{-1}(A))TU_\rho(h)^* = \rho'(A)U_{\rho'}(h)TU_\rho(h)^* = T_h\rho'(A), \end{aligned}$$

thus $T_h \in \text{Hom}(\rho, \rho')$. By assumption, ρ, ρ' have conjugates and thus $\text{Hom}(\rho, \rho')$ is finite dimensional by Proposition 218. Thus $(h, T) \mapsto T_h$ is a finite dimensional representation of the Poincaré group \mathcal{P} . The claim $TU_\rho(h) = U_{\rho'}(h)T$ is equivalent to triviality of this representation. This triviality follows from the non-existence of finite dimensional unitary representations of $\widehat{\mathcal{P}}$ as soon as one produces a positive definite $\widehat{\mathcal{P}}$ -invariant inner product on $\text{Hom}(\rho, \rho')$. For this last step we refer, e.g., to [Roberts, 1990].

- (ii) Apply (i) to $\rho' = \rho$, $U_{\rho'}(h) = \widetilde{U}_\rho(h)$, $T = \text{id}_\rho = \mathbf{1}_{H_0}$ to conclude $U_\rho = \widetilde{U}_\rho$.
- (iii) Using (i) we compute

$$\begin{aligned} \beta_h(T)X_\rho(h) &= (U(h)TU(h)^*)(U(h)U_\rho(h)^*) = U(h)TU_\rho(h^{-1}) \\ &= U(h)U_{\rho'}(h^{-1})T = X_{\rho'}(h)T, \end{aligned}$$

- (iv) See [Roberts, 1990]. ■

NOTE 224. Under weak additional assumptions on the net \mathfrak{A} , it is shown in [Guido and Longo, 1992, Theorem 5.2] that every localized endomorphism of finite dimension is automatically covariant with positive energy! Equivalently, $\Delta_f \subset \Delta_c$, and therefore $\Delta_{fc} = \Delta_f$.

8.5 Statistics in braided tensor $*$ -categories

DEFINITION 225. Let $(\mathcal{C}, \otimes, \mathbf{1})$ be a tensor $*$ -category with unitary braiding $c_{X,Y}$, and suppose that each $X \in \text{Obj}(\mathcal{C})$ has a conjugate. For each $X \in \text{Obj}(\mathcal{C})$, we define the *twist* of X , $\Theta_X \in \text{End}(X)$, by

$$\Theta_X = (\bar{r}^* \otimes \text{id}_X) \circ (\text{id}_{\bar{X}} \otimes c_{X,X}) \circ (\bar{r} \otimes \text{id}_X),$$

where (\bar{X}, r, \bar{r}) is a standard conjugate for X .

NOTE 226. For each $X \in \text{Obj}(\mathcal{C})$, Θ_X is unitary. When X is irreducible, $\text{End}(X) = \mathbb{C}\text{id}_X$, and so $\Theta_X = \omega_X \text{id}_X$, where ω_X is a complex number of unit modulus (called the *statistics phase*). In the case that $c_{X,Y}$ is a symmetry, then $(c_{X,X})^* = c_{X,X}$, and so $(\Theta_X)^* = \Theta_X$. Together with unitarity, this implies that $(\Theta_X)^2 = \text{id}_X$.

DEFINITION 227. Let $(\mathcal{C}, \otimes, \mathbf{1})$ be a tensor $*$ -category with unitary symmetry $c_{X,Y}$. If $X \in \text{Obj}(\mathcal{C})$ is irreducible, we say that X is a *Bosonic object* if $\omega_X = 1$, and we say that X is a *Fermionic object* if $\omega_X = -1$.

NOTE 228. We give a number of justifications for our focus on the category Δ_f of DHR representations with finite dimension.

(i): In the heuristic interpretation of the (unitary equivalence classes) of irreducible DHR representations as corresponding to the types of particles in a QFT, the conjugate $\bar{\rho}$ of a DHR representation ρ corresponds to the antiparticle. It may happen that a particle is its own antiparticle, i.e. $\rho \cong \bar{\rho}$; but the existence of antiparticles seems to be an integral part of relativistic quantum field theories.

(ii): The DHR sectors admitting a conjugate in the above sense are (rough) analogues in the operator algebraic approach to AQFT of Wightman fields with finitely many components. In the Wightman framework [Streater and Wightman, 1964] it is well known that infinite components behave ‘pathologically’ in that the PCT and spin-statistics theorems do not apply to them, and can in fact be violated. In algebraic QFT, these results are reflected in the fact that we cannot even define Bosonic and Fermionic objects that have dimension ∞ , in the sense that they have no conjugates.

(iii): In [Fredenhagen, 1981] it was shown that every massive one-particle representation (cf. Note 207(iii)), which by the mentioned result of [Buchholz and Fredenhagen, 1982] is localizable in space-like cones, has a conjugate in the C^* -tensor category of cone-localizable representations. It therefore seems natural to require existence of conjugates also in the more restrictive setting of double cone localizable representations.

(iv): As pointed out in Note 224, DHR endomorphisms of finite dimension are automatically covariant, provided one accepts the additional conditions on the net \mathfrak{A} needed for this result. Even if one doesn’t wish to appeal to this result, finite dimensionality of the objects is needed (via finite dimensionality of the hom-sets) for the proof of Proposition 223. The latter will be crucial for lifting the Poincaré action from \mathfrak{A} to the field theory \mathfrak{F} in Section 10.

9 FROM FIELDS TO REPRESENTATIONS

In the current section we take the ‘top down’ approach to superselection rules. That is, we are given a field algebra \mathfrak{F} and a gauge group G acting concretely on a Hilbert space \mathcal{H} . We then define the observables as the gauge invariant elements of \mathfrak{F} . The representation of \mathfrak{F} on \mathcal{H} then gives us a preferred set of representations of \mathfrak{A} ; viz. those that can be ‘created from the vacuum representation by the action of local fields.’ Our main mathematical objective in the current section is to show that these representations satisfy the DHR selection criterion. Thus, all superselection sectors that arise in the traditional way — viz. by acting on the vacuum with fields — fall within the purview of DHR superselection theory. (But note: We are restricting attention to *local* fields.)

DEFINITION 229. Let ω_0 be a state on \mathfrak{A} , and let (\mathcal{H}_0, π_0) be the corresponding GNS representation. A *field system with gauge symmetry* for (\mathfrak{A}, ω_0) is a quadruple $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$, where (\mathcal{H}, π) is a representation of \mathfrak{A} , $O \mapsto \mathfrak{F}(O)$ is a net of von Neumann algebras acting irreducibly on \mathcal{H} , G is a strongly compact group of unitary operators on \mathcal{H} , k is a central element of G such that $k^2 = e$, and such that:

- $\alpha)$ (\mathcal{H}_0, π_0) is a subrepresentation of (\mathcal{H}, π) , i.e. there is an isometry $V : \mathcal{H}_0 \rightarrow \mathcal{H}$ such that $V\pi_0 = \pi V$;
- $\beta)$ V maps \mathcal{H}_0 into the subspace of G -invariant vectors of \mathcal{H} ;
- $\gamma)$ the $U \in G$ induce automorphisms that leave each $\mathfrak{F}(O)$ globally fixed, and $\pi(\mathfrak{A}(O))'' \subseteq \mathfrak{F}(O)$ is the set of fixed points under the action of G on $\mathfrak{F}(O)$;
- $\delta)$ for each $O \in \mathcal{K}$, $V(\mathcal{H}_0)$ is cyclic for $\mathfrak{F}(O)$;
- $\varepsilon)$ the fields are local relative to the observables, i.e. $\mathfrak{F}(O_1)$ and $\pi(\mathfrak{A}(O_2))$ commute elementwise whenever O_1 and O_2 are spacelike separated.

A few remarks on the definition of a field system: the fact that \mathfrak{F} is generated by local algebras $\{\mathfrak{F}(O) : O \in \mathcal{K}\}$ means that elements of \mathfrak{F} represent *local* fields — i.e., fields whose excitations can be localized within a bounded spacetime region. Furthermore:

- $\delta)$ is the Reeh-Schlieder Condition: it states that each local region O carries a full set of fields in the sense that these local fields can reach each sector from the vacuum sector. [But note that Condition (δ) only guarantees that sectors in \mathcal{H} can be reached from the vacuum sector. A stronger notion of

completeness would rely on some intrinsic criterion for physical sectors of \mathfrak{A} , and would require that all these sectors be contained in \mathcal{H} ; see Definition 246.]

- γ) can be interpreted as saying that the group G is an internal symmetry group of the field: it does not change the spacetime localization region of field operators.
- ε) is the Relative Locality Condition. Since fields need not be observable, the field algebra is *not* required to satisfy microcausality. However, in the typical situation (i.e. normal commutation relations), field operators localized in one spacetime region either commute or anticommute with field operators localized in a spacelike separated region. Condition (ε) is a weakening of the requirement of normal (Bose/Fermi) commutation relations.

Since G is a compact group of unitary operators acting on \mathcal{H} , we can apply all of the apparatus of the theory of unitary representations of compact groups (see e.g. [Folland, 1995]). In particular, \mathcal{H} decomposes into a direct sum of orthogonal subspaces \mathcal{H}_ξ which reduce the action of G . Thus the reduced unitary representation U_ξ of G on \mathcal{H}_ξ is factorial, i.e. the von Neumann algebra generated by the operators $\{g|_{\mathcal{H}_\xi} : g \in G\}$ is a factor. The representation U_ξ decomposes into a direct sum of *unitarily equivalent* irreducible representations of G . So, there is a privileged direct sum decomposition of \mathcal{H} :

$$\mathcal{H} = \bigoplus_{\xi} \mathcal{H}_\xi,$$

where the subspace \mathcal{H}_ξ is generated by the vectors in \mathcal{H} that transform according to the character ξ (unitary equivalence class of irreducible representations) of G .

In the present section our primary objectives are:

1. Show that the subspaces \mathcal{H}_ξ reduce the action of the observable algebra \mathfrak{A} . So, the representation of \mathfrak{A} on \mathcal{H} decomposes into a direct sum $\bigoplus_{\xi} \pi_\xi$ of representations on the subspaces \mathcal{H}_ξ .
2. Show that each representation $(\mathcal{H}_\xi, \pi_\xi)$ of \mathfrak{A} is factorial, so that the irreducible subrepresentations of $(\mathcal{H}_\xi, \pi_\xi)$ are mutually equivalent. (Hence each character ξ of G labels an equivalence class of irreducible representations of \mathfrak{A} .)
3. Show that $V(\mathcal{H}_0)$ is precisely the subspace of G -invariant vectors in \mathcal{H} . (Hence the character $\mathbf{1}$ of G labels the equivalence class of the vacuum representation of \mathfrak{A} .)
4. Show that each subrepresentation of (\mathcal{H}, π) is a DHR representation. In slogan form, the sectors that can be reached from the vacuum by application of local fields correspond to DHR representations (i.e. representations that are equivalent, modulo any local region, to the vacuum representation).

Regarding objectives (1) and (2), it will suffice to show that $\pi(\mathfrak{A})'' = G'$, because then the von Neumann algebras $\pi(\mathfrak{A})''$ and G'' share the same central projections.

PROPOSITION 230. *If $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ is a field system with gauge symmetry for (\mathfrak{A}, ω_0) then $\pi(\mathfrak{A})' = G''$.*

Our notation will henceforth be simplified if we use g and $U(g)$ ambiguously to denote elements of the unitary group G on \mathcal{H} . That is, we think of $g \rightarrow U(g)$ as the identity representation of G on \mathcal{H} .

Proof. Define $M : \mathfrak{B}(\mathcal{H}) \rightarrow G'$ by

$$M(A) = \int_G U(g)AU(g)^* d\mu(g),$$

where μ is the Haar measure on G . Then M is a faithful, normal projection of norm one from $\mathfrak{B}(\mathcal{H})$ onto G' . Since M is weakly continuous on the unit ball of $\mathfrak{B}(\mathcal{H})$, we have

$$(29) \quad G' = M(\mathfrak{B}(\mathcal{H})) = M(\overline{\mathfrak{F}}) = \overline{M(\mathfrak{F})} = \overline{\pi(\mathfrak{A})}.$$

Thus, $G'' = \pi(\mathfrak{A})'$. ■

It follows then that the factorial subrepresentations of the representation (\mathcal{H}, π) of \mathfrak{A} are in one to one correspondence with the factorial subrepresentations of the action of G on \mathcal{H} .

NOTE 231. Since G is compact each irreducible representation of G is finite dimensional. Let \hat{G} be the set of characters (equivalence class of irreducible representations) of G , and for $\xi \in \hat{G}$, let $d(\xi)$ be the dimension of the underlying Hilbert space. Then the previous result gives a nice intuitive picture of the representation (\mathcal{H}, π) of \mathfrak{A} . For each $\xi \in \hat{G}$, select an irreducible subrepresentation $(\mathcal{H}_\rho, \pi_\rho)$ of the factorial representation $(\mathcal{H}_\xi, \pi_\xi)$. Then we have

$$\pi(A) = \bigoplus_{\xi \in \hat{G}} d(\xi)\pi_\rho(A) = \bigoplus_{\xi \in \hat{G}} (\pi_\rho(A) \otimes I_\rho),$$

where $d(\xi)\pi_\rho(A) = \pi_\rho(A) \oplus \cdots \oplus \pi_\rho(A)$, $d(\xi)$ times, and I_ρ is the identity on an $d(\xi)$ -dimensional Hilbert space.

LEMMA 232. *Let $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ be a field system with gauge symmetry for (\mathfrak{A}, ω_0) . Then \mathcal{H}_0 is separating for \mathfrak{F} .*

Proof. Let $F \in \mathfrak{F}$. If $F\mathcal{H}_0 = \{0\}$ then $\mathcal{E}(F^*F)\mathcal{H}_0 = \{0\}$. Since $\mathcal{E}(F^*F) \in \pi(\mathfrak{A})$ and π_0 is faithful, $\mathcal{E}(F^*F) = 0$. Since \mathcal{E} is faithful, $F = 0$. Therefore, \mathcal{H}_0 is separating for \mathfrak{F} . ■

To obtain further information about the field system $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$, we identify “tensors under the action of G ” in the field algebra \mathfrak{F} . To make sense of this idea, forget momentarily that \mathfrak{F} has a product operation, and consider it merely

as a Banach space. The map $U \mapsto \text{Ad}U$ is a (strongly) continuous representation of the compact group G in $\text{Aut}\mathfrak{F}$, which is of course a subset of the invertible linear operators on \mathfrak{F} . As in the case of a representation of a compact group on a Hilbert space H , the representation of G on \mathfrak{F} decomposes into a direct sum of disjoint representations. An operator $F \in \mathfrak{F}$ is said to *transform according to a representation ρ of G* just in case it is contained in a linear subspace H_ρ of \mathfrak{F} carrying the corresponding representation of G . In fact, we will show that the irreducible subspaces in \mathfrak{F} have a special algebraic property: they have support I .

LEMMA 233. *Let $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ be a field system with gauge symmetry for (\mathfrak{A}, ω_0) , and suppose that \mathfrak{A} satisfies property B relative to ω_0 . Then the net $O \mapsto \pi(\mathfrak{A}(O))''$ satisfies property B.*

Proof. We first establish that $\pi_0|_{\mathfrak{A}(O')}$ is quasiequivalent to $\pi|_{\mathfrak{A}(O')}$ for each double cone O .

By the Relative Locality Condition (ε) , $\mathfrak{F}(O) \subseteq \pi(\mathfrak{A}(O'))'$. By the Reeh-Schlieder Condition (δ) , \mathcal{H}_0 is a cyclic subspace for $\mathfrak{F}(O)$. Thus,

$$\mathcal{H} = [\mathfrak{F}(O)\mathcal{H}_0] \subseteq [\pi(\mathfrak{A}(O'))'\mathcal{H}_0].$$

Let E_0 be the orthogonal projection onto \mathcal{H}_0 . The central support of E_0 in $\pi(\mathfrak{A}(O'))'$ is the projection onto $[\pi(\mathfrak{A}(O'))'E_0(\mathcal{H})]$ [Kadison and Ringrose, 1997, Prop. 5.5.2]. Thus E_0 has central support I in $\pi(\mathfrak{A}(O'))'$, and therefore $(\pi_0|_{\mathfrak{A}(O')}, \mathcal{H}_0)$ and $(\pi|_{\mathfrak{A}(O')}, \mathcal{H})$ are quasiequivalent [Kadison and Ringrose, 1997, Thm. 10.3.3].

Let O_1 be a double cone whose closure is contained in O , and let E be a nonzero projection in $\pi(\mathfrak{A}(O_1))''$. Choose a double cone O_2 that is spacelike separated from O . The preceding argument shows that there is a $*$ -isomorphism φ from $\pi_0(\mathfrak{A}(O_2))''$ to $\pi(\mathfrak{A}(O_2))''$ such that $\varphi(\pi_0(A)) = \pi(A)$ for all $A \in \mathfrak{A}$. This isomorphism φ preserves the net structure: $\varphi[\pi_0(\mathfrak{A}(O_3))] = \pi(\mathfrak{A}(O_3))$ for any double cone O_3 contained in O_2' . Further, since φ is ultraweakly continuous [Kadison and Ringrose, 1997, Cor. 7.1.16], $\varphi[\pi_0(\mathfrak{A}(O_3))''] = \pi(\mathfrak{A}(O_3))''$. In particular, $\varphi(E)$ is a projection in $\pi_0(\mathfrak{A}(O_1))''$. By property B for π_0 , there is an isometry $V \in \pi_0(\mathfrak{A}(O))''$ such that $VV^* = \varphi(E)$. Thus, $W := \varphi^{-1}(V) \in \pi(\mathfrak{A}(O))''$ is an isometry such that $WW^* = E$. Therefore the net $O \mapsto \pi(\mathfrak{A}(O))''$ satisfies property B. ■

DEFINITION 234. Consider the ordered n -tuple (F_1, \dots, F_n) of elements in \mathfrak{F} . We say that this n -tuple *transforms under the action of G according to character ξ* just in case:

1. $F_i^*F_j = 0$ if $i \neq j$; and
2. $\alpha_g(F_i) = \sum_{j=1}^n u_{ij}^\xi(g)F_j$, where u_{ij}^ξ is a set of matrix elements for ξ . That is, for some representation (H, ρ) of G of class ξ , and orthonormal basis $\{e_1, \dots, e_n\}$ for H , $u_{ij}(g) := \langle e_i, \rho(g)e_j \rangle_H$.

NOTE 235. If (F_1, \dots, F_n) is a tensor in \mathfrak{F} transforming according to ξ , then we can always replace the F_i 's with partial isometries V_i with orthogonal ranges. Indeed, let $V_i|F_i|$ be the polar decomposition of F_i . When $i \neq j$, $F_i^*F_j = 0$, and so F_i^* and F_j^* have orthogonal ranges. Recall that if $F = V|F|$, then V annihilates the orthogonal complement of $\text{r}(F^*) = \text{r}(|F|)$ (see [Kadison and Ringrose, 1997, Thm. 6.1.2]). Thus $V_j|F_i| = \delta_{ij}F_j$, and

$$\left(\sum_j u_{ij}^\xi(g)V_j\right)|F_i| = \sum_j u_{ij}^\xi(g)V_j|F_j| = \sum_j u_{ij}^\xi(g)F_j = F_i.$$

By the uniqueness of the polar decomposition, $\sum_j u_{ij}^\xi(g)V_j = V_i$. Hence (V_1, \dots, V_n) is a tensor transforming according to ξ .

DEFINITION 236. Given $\varphi, \psi \in \mathcal{H}_\xi$, define a map $M_{\varphi, \psi}^\xi : \mathfrak{F} \rightarrow \mathfrak{F}$ by

$$M_{\varphi, \psi}^\xi(F) = \int_G \langle \varphi, U(g)\psi \rangle \alpha_g(F) d\mu(g),$$

where μ is the Haar measure on G .

FACT 237. Due to the invariance of μ we have $\alpha_g \circ M_{\varphi, \psi}^\xi(F) = M_{U(g)\varphi, \psi}^\xi(F)$.

LEMMA 238. Let (F_1, \dots, F_n) be a tensor in $\mathfrak{F}(O)$ transforming as a unitary representation of class ξ . Then $F_i(\mathcal{H}_0) \subseteq \mathcal{H}_\xi$, where \mathcal{H}_ξ is the subspace of vectors of \mathcal{H} that transform according to ξ .

Sketch of proof. Let $\varphi \in \mathcal{H}_0$, and let $g \in G$. Then

$$U(g)[M_{\varphi, \psi}^\xi(F)\varphi] = U(g)M_{\varphi, \psi}^\xi(F)U(g)^*\varphi = M_{U(g)\varphi, \psi}^\xi(F)\varphi.$$

Then a straightforward calculation using matrix elements for ξ establishes the result. ■

LEMMA 239. Let ξ be a character of G that occurs nontrivially in the decomposition of the action of G on \mathcal{H} . Then for each double cone O , there is a tensor (F_1, \dots, F_n) in $\mathfrak{F}(O)$ that transforms as a unitary representation of class ξ .

Sketch of proof. Let (ψ_1, \dots, ψ_n) be an orthonormal basis from a G -irreducible subspace of \mathcal{H}_ξ . Let φ be a unit vector in this same subspace. Since \mathcal{H}_0 is cyclic for $\mathfrak{F}(O)$, and $\mathfrak{F}(O)$ is a von Neumann algebra, there is an $F \in \mathfrak{F}(O)$ and a vector $\varphi_0 \in \mathcal{H}_0$ such that $F\varphi_0 = \varphi$. Let $F_i = M_{\psi_i, \varphi}^\xi(F)$. One verifies then that (F_1, \dots, F_n) is the required tensor. ■

LEMMA 240. Let $F_1, \dots, F_n \in \mathfrak{F}(O)$ such that (F_1, \dots, F_n) transforms according to the character ξ . Then if $\overline{O} \subseteq O_1$, there are $X_1, \dots, X_n \in \mathfrak{A}(O_1)$ such that (X_1, \dots, X_n) transforms according to ξ and

$$\sum_{i=1}^n X_i^* X_i = I.$$

Proof. First replace F_1, \dots, F_n with partial isometries V_1, \dots, V_n , then let $V = \sum_{i=1}^n V_i$. Since the V_i have orthogonal ranges, V is a partial isometry, and $V^*V = \sum_{i=1}^n V_i^*V_i$. A straightforward calculation shows that $\alpha_g(V^*V) = V^*V$ for all $g \in G$. Thus, $E = V^*V$ is a projection in $\mathfrak{F}(O) \cap G' = \pi(\mathfrak{A}(O))''$. By Lemma 233, $O \mapsto \pi(\mathfrak{A}(O))''$ satisfies property B. Thus, there is an isometry $W \in \pi(\mathfrak{A}(O_1))''$ with $WW^* = E$. For $i = 1, \dots, n$, let $X_i = V_iW$. Then the tensor (X_1, \dots, X_n) transforms according to ξ , and

$$\sum_{i=1}^n X_i^*X_i = W^* \left(\sum_{i=1}^n V_i^*V_i \right) W = I.$$

■

LEMMA 241. *Let $\mathcal{H}_\xi \subseteq \mathcal{H}$ be the range of a central projection in $\pi(\mathfrak{A})''$. Then for each double cone O , \mathcal{H}_ξ is cyclic for $\mathfrak{F}(O)$.*

Proof. Let O_1 be a double cone such that $\overline{O_1} \subseteq O$. By the Reeh-Schlieder Condition, \mathcal{H}_0 is cyclic for $\mathfrak{F}(O_1)$. By Lemma 239, there is a tensor (F_1, \dots, F_n) in $\mathfrak{F}(O_1)$ that transforms according to the representation (H, ρ) of G . By Lemma 240, there is a tensor (X_1, \dots, X_n) in $\mathfrak{F}(O)$ that transforms that same way, and such that $\sum_{i=1}^n X_i^*X_i = I$. Then

$$\mathfrak{F}(O)\mathcal{H}_0 = \mathfrak{F}(O) \sum_{i=1}^n X_i^*X_i\mathcal{H}_0 \subseteq \mathfrak{F}(O)\mathcal{H}_\xi,$$

where the final inclusion follows by Lemma 238. Therefore \mathcal{H}_ξ is cyclic for $\mathfrak{F}(O)$. ■

DEFINITION 242. Let $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$ be the category of subrepresentations of the representation (\mathcal{H}, π) of \mathfrak{A} . We mean to take $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$ as a full subcategory of the category of all representations of \mathfrak{A} , i.e. the hom-sets between representations in $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$ are the same as the hom-sets in the larger category.

PROPOSITION 243. *Let $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ be a field system with gauge symmetry for (\mathfrak{A}, ω_0) . Then there is a faithful functor $F : \text{Rep}_{\mathfrak{F}}\mathfrak{A} \rightarrow \text{DHR}(\mathfrak{A})$.*

Proof. Suppose that (\mathcal{H}', π') is an object of $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$. That is, there is an isometry $V : \mathcal{H}' \rightarrow \mathcal{H}$ such that $V\pi' = \pi V$. We subsequently identify \mathcal{H}' with its image in \mathcal{H} , and treat π' as mapping into $\mathfrak{B}(\mathcal{H})$. We must show that (\mathcal{H}', π') is in $\text{DHR}(\mathfrak{A})$; that is, for any double cone O , $(\mathcal{H}', \pi'|_{\mathfrak{A}(O)})$ is unitarily equivalent to $(\mathcal{H}_0, \pi_0|_{\mathfrak{A}(O)})$.

Let $\bar{\pi} = \pi|_{\mathfrak{A}(O')}$. Since $E_L, E_\xi \in \pi(\mathfrak{A})' \subseteq \pi(\mathfrak{A}(O'))'$, E_L and E_ξ reduce $\bar{\pi}$. We first establish that E_L and E_ξ have the same central support in $\pi(\mathfrak{A}(O'))^-$, from which it follows that $E_L\bar{\pi}$ and $E_\xi\bar{\pi}$ are quasiequivalent [Kadison and Ringrose, 1997, Thm. 10.3.3].

By the Relative Locality Condition (ε), $\mathfrak{F}(O) \subseteq \pi(\mathfrak{A}(O'))'$. By the Reeh-Schlieder Condition (δ), $E_l\mathcal{H}$ is a cyclic subspace for $\mathfrak{F}(O)$. Thus,

$$\mathcal{H} = [\mathfrak{F}(O)E_l(\mathcal{H})] \subseteq [\pi(\mathfrak{A}(O'))'E_l(\mathcal{H})].$$

Similarly, Lemma 241 entails that $E_\xi\mathcal{H}$ is a cyclic subspace for $\mathfrak{F}(O)$, and so $[\pi(\mathfrak{A}(O'))'E_\xi(\mathcal{H})] = \mathcal{H}$. However, the central support of E_l in $\pi(\mathfrak{A}(O'))'$ is the projection onto $[\pi(\mathfrak{A}(O'))'E_0(\mathcal{H})]$, and similarly for E_ξ [Kadison and Ringrose, 1997, Prop. 5.5.2]. Thus, E_l and E_ξ have central support I in $\pi(\mathfrak{A}(O'))'$. Therefore, $(\pi_0|_{\mathfrak{A}(O')}, \mathcal{H}_0)$ and $(\pi_\xi|_{\mathfrak{A}(O')}, \mathcal{H}_\xi)$ are quasiequivalent, i.e. there is a $*$ -isomorphism $\varphi : \pi_0(\mathfrak{A}(O')) \rightarrow \pi_\xi(\mathfrak{A}(O'))$ such that $\varphi(\pi_0(A)) = \pi_\xi(A)$ for all $A \in \mathfrak{A}(O')$.

The previous reasoning also shows (by replacing O with a spacelike separated double cone) that for each double cone O , $(\pi_0|_{\mathfrak{A}(O)}, \mathcal{H}_0)$ is quasiequivalent to $(\pi_\xi|_{\mathfrak{A}(O)}, \mathcal{H}_\xi)$. Thus, in particular, since the net $O \rightarrow \pi_0(\mathfrak{A}(O))''$ of von Neumann algebras satisfies property B (by assumption), so does the net $O \rightarrow \pi_\xi(\mathfrak{A}(O))''$.

To establish that $(\pi_0|_{\mathfrak{A}(O')}, \mathcal{H}_0)$ and $(\pi_\xi|_{\mathfrak{A}(O')}, \mathcal{H}_\xi)$ are unitarily equivalent, we will use the following result ([Kadison and Ringrose, 1997, Theorem 7.2.9]):

Let $\mathfrak{R}_j, j = 1, 2$, be von Neumann algebras acting on Hilbert spaces \mathcal{H}_j respectively. Suppose that for $j = 1, 2$, there is a vector $x_j \in \mathcal{H}_j$ that is cyclic and separating for \mathfrak{R}_j . If $\alpha : \mathfrak{R}_1 \rightarrow \mathfrak{R}_2$ is a $$ isomorphism then there is a unitary operator U from \mathcal{H}_1 to \mathcal{H}_2 that implements α .*

in conjunction with the fact ([Kadison and Ringrose, 1997, Exercise 9.6.32]):

If \mathfrak{R} is a von Neumann algebra acting on a separable Hilbert space \mathcal{H} , and if \mathfrak{R}' is properly infinite, then there is vector $x \in \mathcal{H}$ that is cyclic and separating for \mathfrak{R} .

By Proposition 71, $\pi_0(\mathfrak{A}(O'))'$ and $\pi_\xi(\mathfrak{A}(O'))'$ are properly infinite. By assumption, \mathcal{H}_0 is separable. Thus, it will suffice to show that \mathcal{H}_ξ is separable. Since π_ξ is irreducible, each non-zero vector $x \in \mathcal{H}_\xi$ is cyclic for $\pi_\xi(\mathfrak{A})$. Thus, \mathcal{H}_ξ is the closure of the union of $\pi_\xi(\mathfrak{A}(O_n))x$ for an increasing sequence O_n of double cones. Hence it suffices to show that $\pi_\xi(\mathfrak{A}(O))x$ is separable for each $O \in \mathcal{K}$. Since \mathcal{H}_0 is separable, the unit ball of $\mathfrak{B}(\mathcal{H}_0)$ is compact metrizable [Kadison and Ringrose, 1997, Thm. 5.1.3; Exercise 5.7.7]. Since the unit ball of $\pi_0(\mathfrak{A}(O))''$ is a closed subset of the unit ball of $\mathfrak{B}(\mathcal{H}_0)$, it is also compact metrizable. But $\pi_\xi(\mathfrak{A}(O))''$ is $*$ isomorphic, hence ultraweakly homeomorphic, to $\pi_0(\mathfrak{A}(O))''$. Therefore, the unit ball of $\pi_\xi(\mathfrak{A}(O))''$ is compact metrizable, hence separable, in the weak operator topology. It follows that $\pi_\xi(\mathfrak{A}(O))''x$ is separable. ■

In Proposition 206 it was shown that there is a faithful, essentially surjective functor F' from the category $DHR(\mathfrak{A})$ of DHR representations to the category Δ of localized transportable morphisms of \mathfrak{A} . So, the previous Proposition entails that $F' \circ F$ is a faithful functor from $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$ into Δ . We subsequently replace $F' \circ F$ with just F .

Recall that Δ_f is the full subcategory of Δ of objects with conjugates. The final thing we need to show in this Section is that the image of each object in $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$ under F is isomorphic to an object in Δ_f . That is, we need to show that the image object has a conjugate.

Sketch of proof. One shows that the subrepresentations of G on \mathcal{H} are closed under taking conjugates. This can be proven by noting a correspondence between the action of G on \mathcal{H} and the action of G on \mathfrak{F} . Then use the fact that \mathfrak{F} is a $*$ -algebra. Thus, for each irreducible subrepresentation π_ρ of π , there is an irreducible subrepresentation $\pi_{\bar{\rho}}$ of π . Verify that $(F' \circ F)(\pi_{\bar{\rho}})$ is a conjugate for $(F' \circ F)(\pi_\rho)$. ■

Therefore, $F' \circ F$ is a faithful functor from $\text{Rep}_{\mathfrak{F}}\mathfrak{A}$ into Δ_f . So we have shown:

Each representation of \mathfrak{A} that arises from its being taken as the gauge invariant part of a field algebra is a representation of the form $\pi_0 \circ \rho$ with $\rho \in \text{Obj}(\Delta_f)$.

Thus, the study of Δ_f encompasses the study of representations that arise from the approach that begins with a field algebra.

We said above that in the “normal” situation, field operators in $\mathfrak{F}(O_1)$ with either commute or anticommute with field operators in $\mathfrak{F}(O_2)$ when O_1 and O_2 are spacelike separated. To be more precise, we would expect for a Bose field operator to commute both with other Bose field operators, as well as with Fermi field operators; and we would expect for a pair of Fermi field operators to anticommute. But what are Bose and Fermi field operators? The distinction between the two is defined in terms of the privileged element k of the gauge group G .

DEFINITION 244. If $\alpha_k(F) = F$ then F is said to be a *Bose field operator*; and if $\alpha_k(F) = -F$ then F is said to be a *Fermi field operator*.

We define a Bosonic sector in \mathcal{H} to be a subspace \mathcal{H}_ξ such that $U(k)|_{\mathcal{H}_\xi} = I$, and a Fermionic sector in \mathcal{H} to be a subspace \mathcal{H}_ξ such that $U(k)|_{\mathcal{H}_\xi} = -I$. It then follows that Bosonic field operators create Bosonic sectors from the vacuum, and Fermionic field operators create Fermionic sectors from the vacuum.

We can now make sense of the notion of normal commutation relations: Bose field operators should commute with each other and commute with Fermionic field operators. Fermionic field operators should anticommute with each other.

DEFINITION 245. A local operator algebra system of fields $(\pi, (G, k), \mathfrak{F})$ is said to *satisfy normal commutation relations* just in case the local field algebras satisfy *graded local commutativity*: If O_1 and O_2 are spacelike, and $F_\sigma \in \mathfrak{F}(O_1)$, $F'_\sigma \in \mathfrak{F}(O_2)$ are such that $\alpha_k(F_\sigma) = \sigma F_\sigma$ and $\alpha_k(F'_\sigma) = \sigma F'_\sigma$, ($\sigma = \pm$), then

$$F_+F'_+ = F'_+F_+, \quad F_+F'_- = F'_-F_+, \quad F_-F'_- = -F'_-F_-.$$

10 FROM REPRESENTATIONS TO FIELDS

The preceding section derives properties of representations of \mathfrak{A} , given that these representations are created from the vacuum representation by the action of local fields on the vacuum. But such an approach will seem at best heuristic to the Algebraic Imperialist. From the Imperialist’s point of view, the entire content of the theory is contained in the abstract net \mathfrak{A} of observable algebras.

On the one hand, the Imperialist might be an eliminativist about fields and gauge group. On the other hand, the Imperialist might claim that the fields and gauge group are physically significant, but only because they can be ‘reconstructed’ from the net of observable algebras. In order to justify this latter stance, the Imperialist would need to accomplish the following:

Task: Try to reconstruct, in a mathematically rigorous fashion, the entire apparatus of QFT — fields, gauge groups, etc. — from the net of observable algebras.

A quixotic task indeed! For one, philosophers seemed to have settled that theory is always underdetermined by data; and so we should not expect to be able to find the full theoretical apparatus hidden within the net of observable algebras. But there is a surprise in store: the Task was undertaken, and was achieved. The DR Reconstruction Theorem shows in a fully rigorous and precise way that the DHR category encodes all the information needed to reconstruct, uniquely, the fields and the gauge group. This section provides the details of the reconstruction.

DEFINITION 246. A field system with gauge symmetry $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ for (\mathfrak{A}, ω_0) is said to be *complete* if the representation π of \mathfrak{A} contains copies of all representations in the DHR category $DHR(\mathfrak{A})$ of \mathfrak{A} .

DEFINITION 247. Two field systems with gauge symmetry $(\pi_1, \mathcal{H}_1, \mathfrak{F}_1, G_1)$ and $(\pi_2, \mathcal{H}_2, \mathfrak{F}_2, G_2)$ for (\mathfrak{A}, ω_0) are said to be *equivalent* if there exists a unitary operator $W : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ such that:

$$\begin{aligned} W\pi_1(A) &= \pi_2(A)W, \quad \forall A \in \mathfrak{A}, \\ WU(G_1) &= U(G_2)W, \\ W\mathfrak{F}_1(O) &= \mathfrak{F}_2(O)W, \text{ for each double cone } O. \end{aligned}$$

Doplicher-Roberts Reconstruction Theorem. Let (\mathfrak{A}, ω_0) be a net of observable algebras satisfying duality and property B relative to a privileged ‘vacuum’ state ω_0 . Then there exists a field system with gauge symmetry $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ for (\mathfrak{A}, ω_0) that is complete, and that has normal commutation relations. Any complete, normal field system for (\mathfrak{A}, ω_0) is equivalent to $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$.

The proof of the reconstruction theorem is contained in [Doplicher and Roberts, 1989] and [Doplicher and Roberts, 1990]. In this article, we give an alternative proof, based on Deligne’s embedding theorem [Deligne, 1990], and results obtained by Roberts [Roberts, ND] prior to obtaining the full proof of the reconstruction theorem.

In outline, the theorem shows first — as was essentially established in [Doplicher *et al.*, 1971] — that the DHR superselection sectors naturally have the structure of a braided tensor $*$ -category with conjugates — and when the spacetime dimension is three or greater, we can replace “braided” with “symmetric.” Now, until the late 1980’s, this first result was merely suggestive: it is known that the category $\text{Rep}_f G$ of representations of a compact group G on finite dimensional (super) Hilbert spaces is a symmetric tensor $*$ -category with conjugates. Hence, the category of DHR superselection sectors seems to have all the structure of $\text{Rep}_f G$ for some compact G . By the classical Tannaka-Krein duality theorem, it is possible to reconstruct G from $\text{Rep}_f G$. Furthermore, Roberts [Roberts, ND] proved the conditional claim that *if* the category of superselection sectors was equivalent to the category $\text{Rep}_f G$ for some compact G , then the field algebra \mathfrak{F} could be reconstructed.

But there is a crucial difference between the category of superselection sectors and the category $\text{Rep}_f G$. The category $\text{Rep}_f G$ is “concrete” — it comes with an embedding into the category of Hilbert spaces, namely the forgetful functor, and hence its objects can be regarded as structured sets. It is also precisely the existence of such an embedding that is needed to construct a field algebra, because one needs the objects in the category to have “internal structure,” as, for example, an object in the category \mathcal{H} of Hilbert spaces is a structured set. Before we state the embedding theorem, whose proof is given in the Appendix, we need some preparatory definitions concerning ‘supermathematics’.

10.1 Supermathematics and the embedding theorem

DEFINITION 248. A *super vector space*, alternatively a \mathbb{Z}_2 -graded vector space, is a vector space V with a distinguished decomposition $V = V_+ \oplus V_-$. The subspace V_+ is called the *even* subspace, and V_- is called the *odd* subspace. Elements of $V_+ \cup V_- =: V_h$ are called *homogeneous*. Define the parity function ω on the

homogeneous elements by setting $\omega(v) = \pm 1$ if $v \in V_{\pm}$. A morphism between two super vector spaces is a linear mapping $T : V \rightarrow W$ such that $T(V_{\pm}) \subseteq W_{\pm}$. We let SVect denote the category of super vector spaces. A *super Hilbert space* is a super vector space V with a positive definite inner product such that $V_- \perp V_+$. We use \mathcal{SH} to denote the category of super Hilbert spaces.

We now define operations that make SVect into a symmetric tensor category. It is straightforward to verify that the set $\text{Mor}(V, W)$ of morphisms between two super vector spaces is a linear subspace of $B(V, W)$. Thus, SVect is a linear category.

If V and W are super vector spaces, then their direct sum is the vector space $V \oplus W$ with even subspace $V_+ \oplus W_+$ and odd subspace $V_- \oplus W_-$. We define the monoidal product in SVect as the vector space $V \otimes W$ whose even and odd subspaces are defined by

$$(V \otimes W)_{\sigma} = \bigoplus_{\sigma' \sigma'' = \sigma} V_{\sigma'} \otimes W_{\sigma''},$$

where $\sigma = \pm$. Thus,

$$\begin{aligned} (V \otimes W)_+ &= (V_+ \otimes W_+) \oplus (V_- \otimes W_-), \\ (V \otimes W)_- &= (V_+ \otimes W_-) \oplus (V_- \otimes W_+). \end{aligned}$$

The monoidal unit is \mathbb{C} , with even subspace \mathbb{C} .

DEFINITION 249. For two super vector spaces V, W , we define the symmetry isomorphism

$$c_{V,W} : V \otimes W \rightarrow W \otimes V,$$

by setting

$$c_{V,W}(v \otimes w) = (-1)^{(1-\omega(v))(1-\omega(w))/4} w \otimes v, \quad \forall v \in V_h, \forall w \in W_h.$$

on homogeneous simple tensors, and then by extending linearly.

PROPOSITION 250. *Both $(\text{SVect}, \otimes, \mathbb{C}, c_{V,W})$ and $(\mathcal{SH}, \otimes, \mathbb{C}, c_{V,W})$ are symmetric tensor $*$ -categories.*

NOTE 251. By the coherence theorem \mathcal{SH} is equivalent to a strict symmetric tensor $*$ -category, which we will also denote by \mathcal{SH} .

DEFINITION 252. A *supergroup* is a pair (G, k) where G is a group and k is a central element in G such that $k \cdot k = e$. A morphism between two supergroups (G_1, k_1) and (G_2, k_2) is a group homomorphism $\varphi : G_1 \rightarrow G_2$ such that $\varphi(k_1) = k_2$.

DEFINITION 253. A *(unitary) representation* π of a supergroup (G, k) is a super Hilbert space $V = V_+ \oplus V_-$ together with a (unitary) representation π of G on V such that $\pi(k)|_{V_{\pm}} = \pm \text{id}_{V_{\pm}}$. The representations $\text{Rep}(G, k)$ of (G, k) form a symmetric tensor $*$ -category with tensor product and symmetry inherited from \mathcal{SH} , and monoidal unit the trivial representation of (G, k) on \mathbb{C} .

NOTE 254. Let \mathcal{SH}_f be the full subcategory of finite dimensional super Hilbert spaces. For a supergroup (G, k) we denote by $\text{Rep}_f(G, k)$ the full subcategory of finite dimensional representations of (G, k) . The categories \mathcal{SH}_f and $\text{Rep}_f(G, k)$ are semisimple and have conjugates (see the Appendix for more on this terminology). Also, there is a canonical forgetful functor $K : \text{Rep}_f(G, k) \rightarrow \mathcal{SH}_f$.

We now move on to the statement of the Embedding Theorem, which will be required for the reconstruction of the field algebra and gauge group. For more on supermathematics, we refer the reader to [Varadarajan, 2004; Deligne and Morgan, 1999]. (But note that DHR superselection theory is not concerned with supersymmetry in the sense of a symmetry transforming Bosonic and Fermionic fields into each other. Also, our definition of a supergroup is idiosyncratic.)

The Embedding Theorem. Let \mathcal{SH}_f be the category of finite-dimensional super Hilbert spaces over \mathbb{C} . Let $(\mathcal{C}, \otimes, \mathbf{1}, c_{X,Y})$ be a tensor C^* -category with unitary symmetry $c_{X,Y}$, conjugates, direct sums, sub-objects, and irreducible monoidal unit $\mathbf{1}$. (Such a category is called an *STC** in the Appendix.) Then

1. There is a faithful symmetric tensor $*$ -functor $E : \mathcal{C} \rightarrow \mathcal{SH}_f$.
2. There is a compact supergroup (G, k) , where G is the group of unitary natural monoidal transformations of E , and an equivalence $F : \mathcal{C} \rightarrow \text{Rep}_f(G, k)$ of symmetric tensor $*$ -categories such that $E = F \circ K$, where $K : \text{Rep}_f(G, k) \rightarrow \mathcal{SH}_f$ is the forgetful functor.

NOTE 255. The embedding theorem is proven in Appendix B. In its proof we assume the tensor category \mathcal{C} to be strict and we will work with the strictification \mathcal{SH} of the category of super Hilbert spaces. In view of the coherence theorem for symmetric tensor categories the strictness assumptions do not limit the generality of the result. The tensor functor $F : \mathcal{C} \rightarrow \mathcal{SH}_f$ that we construct will, however, not be a strict tensor functor. In the construction of the field net below we do pretend for notational simplicity that F is strict. We will comment on this issue again at the end of this section.

10.2 Construction of the field net, algebraic aspects

We now apply the Embedding Theorem to the case of the DHR category Δ_f of localized transportable morphisms with finite dimension. In particular, we show that given an embedding $E : \Delta_f \rightarrow \mathcal{SH}_f$, it is possible to construct a local system of field algebras $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$. This strategy of reconstruction is based on the unpublished manuscript [Roberts, ND], which assumes the existence of an embedding (or fiber) functor. The actual existence theorem for the embedding

functor — which is based on the work of Tannaka and Deligne, but incorporates more recent simplifications — can be found in the Appendix.

DEFINITION 256. As a set, the field algebra \mathfrak{F}_0 consists of equivalence classes of triples (A, ρ, ψ) , with $A \in \mathfrak{A}$, $\rho \in \text{Obj}(\Delta_f)$, and $\psi \in E(\rho)$, modulo the equivalence relation

$$(AT, \rho, \psi) = (A, \rho', E(T)\psi),$$

for $T \in \text{Hom}(\rho, \rho')$. Since $E(\text{id}_\rho) = \text{id}_{E(\rho)}$ we have $(\lambda A, \rho, \psi) = (A, \rho, \lambda\psi)$. Subsequently, we do not distinguish notationally between a triple (A, ρ, ψ) and its equivalence class.

PROPOSITION 257. \mathfrak{F}_0 is a complex vector space under the operations:

$$(30) \quad \lambda(A, \rho, \psi) := (\lambda A, \rho, \psi), \quad \lambda \in \mathbb{C},$$

and

$$(31) \quad (A_1, \rho_1, \psi_1) + (A_2, \rho_2, \psi_2) := (A_1 W_1^* + A_2 W_2^*, \rho, E(W_1)\psi_1 + E(W_2)\psi_2),$$

where $\psi_i \in E(\rho_i)$ and $W_i \in \text{Hom}(\rho_i, \rho)$ are isometries with

$$(32) \quad W_1 W_1^* + W_2 W_2^* = \text{id}_\rho.$$

In addition,

$$\begin{aligned} (A_1, \rho, \psi) + (A_2, \rho, \psi) &= (A_1 + A_2, \rho, \psi) \\ (A, \rho, \psi_1) + (A, \rho, \psi_2) &= (A, \rho, \psi_1 + \psi_2). \end{aligned}$$

Therefore, identifying \mathfrak{A} with $\{(A, \iota, 1) : A \in \mathfrak{A}, 1 \in \mathbb{C} \equiv E(\iota)\}$, \mathfrak{A} becomes a linear subspace of \mathfrak{F}_0 ; and identifying $E(\rho)$ with $\{(I, \rho, \psi) : \psi \in E(\rho)\}$, $E(\rho)$ becomes a linear subspace of \mathfrak{F}_0 .

Proof. We first verify that the operations are well defined. Scalar multiplication is well defined since for any $T \in \text{Hom}(\rho, \rho')$, $(\lambda A, \rho', E(T)\psi) = ((\lambda A)T, \rho, \psi) = (\lambda(AT), \rho, \psi)$. To show that addition is well defined, we first establish that Eqn. (31) is independent of the choice of W_1 and W_2 . If $W'_i \in \text{Hom}(\rho_i, \rho')$ is another such choice then

$$\begin{aligned} & \left(A_1 W_1'^* + A_2 W_2'^*, \rho', E(W_1')\psi_1 + E(W_2')\psi_2 \right) \\ &= \left(A_1 W_1^* + A_2 W_2^* (W_1 W_1'^* + W_2 W_2'^*), \rho', E(W_1')\psi_1 + E(W_2')\psi_2 \right) \\ &= \left(A_1 W_1^* + A_2 W_2^*, \rho, E(W_1 W_1'^* + W_2 W_2'^*) E(W_1')\psi_1 + E(W_2')\psi_2 \right) \\ &= \left(A_1 W_1^* + A_2 W_2^*, \rho, E(W_1)\psi_1 + E(W_2)\psi_2 \right). \end{aligned}$$

To see that addition is independent of equivalence classes, let $T_i \in \text{Hom}(\rho_i, \rho'_i)$, W_i isometries in $\text{Hom}(\rho_i, \rho)$, and W'_i isometries in $\text{Hom}(\rho'_i, \rho)$. Then,

$$\begin{aligned} & (A_1 T_1, \rho_1, \psi_1) + (A_2 T_2, \rho_2, \psi_2) \\ &= \left(A_1 T_1 W_1^* + A_2 T_2 W_2^*, \rho, E(W_1)\psi_1 + E(W_2)\psi_2 \right) \\ &= \left((A_1 W_1'^* + A_2 W_2'^*)(W_1' T_1 W_1^* + W_2' T_2 W_2^*), \rho, E(W_1)\psi_1 + E(W_2)\psi_2 \right) \\ &= \left(A_1 W_1'^* + A_2 W_2'^*, \rho, E(W_1')E(T_1)\psi_1 + E(W_2')E(T_2)\psi_2 \right) \\ &= (A_1, \rho'_1, E(T_1)\psi_1) + (A_2, \rho'_2, E(T_2)\psi_2). \end{aligned}$$

To prove additivity in the first argument, choose $\sigma = \rho \oplus \rho$, and $W_i \in \text{Hom}(\rho, \sigma)$ the corresponding isometries. Then

$$\begin{aligned} (A_1, \rho, \psi) + (A_2, \rho, \psi) &= (A_1 W_1^* + A_2 W_2^*, \sigma, (E(W_1) + E(W_2))\psi) \\ &= (A_1 W_1^* + A_2 W_2^*, \sigma, E(W_1 + W_2)\psi) \\ &= ((A_1 W_1^* + A_2 W_2^*)(W_1 + W_2), \rho, \psi) \\ &= (A_1 + A_2, \rho, \psi). \end{aligned}$$

Finally, to prove additivity in the second argument, choose $\sigma = \rho \oplus \rho$, and $W_i \in \text{Hom}(\rho, \sigma)$ the corresponding isometries. Then

$$\begin{aligned} (A, \rho, \psi_1) + (A, \rho, \psi_2) &= (A W_1^* + A W_2^*, \sigma, E(W_1)\psi_1 + E(W_2)\psi_2) \\ &= (A(W_1^* + W_2^*), \sigma, E(W_1)\psi_1 + E(W_2)\psi_2) \\ &= (A, \rho, E(W_1^* + W_2^*)(E(W_1)\psi_1 + E(W_2)\psi_2)) \\ &= (A, \rho, \psi_1 + \psi_2). \end{aligned}$$

■

PROPOSITION 258. *The complex linear space \mathfrak{F}_0 becomes an algebra if we define*

$$(33) \quad (A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2) := \left(A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2 \right),$$

where $\psi_i \in E(\rho_i)$, $i = 1, 2$. Furthermore, \mathfrak{A} is a subalgebra of \mathfrak{F}_0 , and the equivalence class of $(I, \iota, 1)$ is a multiplicative identity, where I is the multiplicative identity of \mathfrak{A} , and $1 \in E(\iota) = \mathbb{C}$.

Proof. We first verify that Eqn. (33) is well-defined on \mathfrak{F}_0 . Let $T_i \in \text{Hom}(\rho_i, \rho'_i)$. Recalling that $T_1 \times T_2 = \rho'_1(T_2)T_1$, we have

$$\begin{aligned} (A_1 T_1, \rho_1, \psi_1)(A_2 T_2, \rho_2, \psi_2) &= (A_1 T_1 \rho_1(A_2 T_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2) \\ &= (A_1 \rho'_1(A_2 T_2) T_1, \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2) \\ &= (A_1 \rho'_1(A_2) \rho'_1(T_2) T_1, \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2) \\ &= (A_1 \rho'_1(A_2)(T_1 \times T_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2) \\ &= (A_1 \rho'_1(A_2), \rho'_1 \otimes \rho'_2, E(T_1 \times T_2)(\psi_1 \otimes \psi_2)) \\ &= (A_1 \rho'_1(A_2), \rho'_1 \otimes \rho'_2, E(T_1)\psi_1 \otimes E(T_2)\psi_2) \\ &= (A_1, \rho'_1, E(T_1)\psi_1)(A_2, \rho'_2, E(T_2)\psi_2). \end{aligned}$$

A straightforward calculation shows that multiplication is associative. For distributivity, let $W_i \in \text{Hom}(\rho_i, \rho)$. Then,

$$\begin{aligned} & \left[(A_1, \rho_1, \psi_1) + (A_2, \rho_2, \psi_2) \right] (A_3, \rho_3, \psi_3) \\ &= ((A_1 W_1^* + A_2 W_2^*) \rho(A_3), \rho \otimes \rho_3, (E(W_1)\psi_1 + E(W_2)\psi_2) \otimes \psi_3) \\ &= (A_1 \rho_1(A_3) W_1^* + A_2 \rho_2(A_3) W_2^*, \rho \otimes \rho_3, (E(W_1)\psi_1) \otimes \psi_3 + (E(W_2)\psi_2) \otimes \psi_3) \\ &= (A_1 \rho_1(A_3)(W_1^* \times I_{\rho_3}) + A_2 \rho_2(A_3)(W_2^* \times I_{\rho_3}), \rho \otimes \rho_3, (E(W_1)\psi_1) \otimes \psi_3 + \\ & \quad + (E(W_2)\psi_2) \otimes \psi_3) \\ &= (A_1, \rho_1, \psi_1)(A_3, \rho_3, \psi_3) + (A_2, \rho_2, \psi_2)(A_3, \rho_3, \psi_3). \end{aligned}$$

■

We will need the following basic lemma from linear algebra.

DEFINITION 259. If H, H' are Hilbert spaces and $S \in \text{Hom}(H \otimes H', \mathbb{C})$, then we define an antilinear mapping $\mathcal{J}S : H \rightarrow H'$ by setting

$$((\mathcal{J}S)x, x') = S(x \otimes x'), \quad \forall x \in H, \forall x' \in H'.$$

LEMMA 260.

1. \mathcal{J} is antilinear: $\mathcal{J}(\lambda S) = \bar{\lambda}(\mathcal{J}S)$, and $\mathcal{J}(S_1 + S_2) = \mathcal{J}S_1 + \mathcal{J}S_2$.
2. If $T \in \text{Hom}(H', H)$ then

$$\begin{aligned} T \circ (\mathcal{J}S) &= \mathcal{J}(S \circ (I_H \otimes T^*)), \\ (\mathcal{J}S) \circ T &= \mathcal{J}(S \circ (T \otimes I_{H'})). \end{aligned}$$

3. If $S' \in \text{Hom}(H' \otimes H'', \mathbb{C})$ then $(\mathcal{J}S') \circ (\mathcal{J}S) = (S \otimes I_{H''}) \circ (I_H \otimes S'^*)$.
4. Let $S_1 \in \text{End}(H_1 \otimes H'_1, \mathbb{C})$ and $S_2 \in \text{End}(H_2 \otimes H'_2, \mathbb{C})$. Then

$$(\mathcal{J}S_2 \otimes \mathcal{J}S_1) \circ \Sigma_{H_1, H_2} = \mathcal{J}[S_1 \circ (I_{H_1} \otimes S_2 \otimes I_{H'_1})].$$

Proof. Straightforward. A nice exercise in basic linear algebra. ■

NOTE 261. We will apply the previous Lemma to super Hilbert spaces. But we will take $\Sigma_{H, H'}$ to be the ordinary symmetry on the category \mathcal{H}_f of finite dimensional Hilbert spaces.

LEMMA 262. Let $T \in \text{Hom}(\rho, \rho')$ and pick solutions $(\bar{\rho}, R, \bar{R})$ and $(\bar{\rho}', R', \bar{R}')$ of the conjugate equations with respect to ρ and ρ' ; that is, $R \in \text{Hom}(\iota, \bar{\rho} \otimes \rho)$, $\bar{R} \in \text{Hom}(\iota, \rho \otimes \bar{\rho})$ such that $(\bar{R}^* \times I_\rho) \circ (I_\rho \times R) = I_\rho$, $(R^* \times I_{\bar{\rho}}) \circ (I_{\bar{\rho}} \times \bar{R}) = I_{\bar{\rho}}$, and analogously for R' and \bar{R}' . Set

$$\bar{T} := (I_{\bar{\rho}} \times \bar{R}'^*) \circ (I_{\bar{\rho}} \times T \times 1_{\bar{\rho}'}) \circ (R \times I_{\bar{\rho}'}) = \bar{\rho}(\bar{R}'^* T)R.$$

Then $\bar{T} \in \text{Hom}(\bar{\rho}', \bar{\rho})$ and

$$(34) \quad (I_{\bar{\rho}} \times T) \circ R = (\bar{T} \times I_{\rho'}) \circ R',$$

$$(35) \quad (I_{\rho} \times \bar{T}^*) \circ \bar{R} = (T^* \times I_{\bar{\rho}'}) \circ \bar{R}'.$$

Proof. For Eqn. (34), we have

$$\begin{aligned} (\bar{T} \times I_{\rho'}) \circ R' &= \bar{T}R' = \bar{\rho}(\bar{R}'^*T)RR' = \bar{\rho}(\bar{R}'^*T\rho(R'))R \\ &= \bar{\rho}(\bar{R}'^*\rho'(R')T)R, \end{aligned}$$

where we used the definition of \times for the first equality, the definition of \bar{T} for the second equality, $R \in \text{Hom}(\iota, \bar{\rho} \otimes \rho)$ for the third equality, and $T \in \text{Hom}(\rho, \rho')$ for the fourth equality. But by the conjugate equations, $\bar{R}'^*\rho'(R') = (\bar{R}'^* \times I_{\rho'}) \circ (I_{\rho'} \times R') = I_{\rho'} = I$, and hence $(\bar{T} \times I_{\rho'}) \circ R' = \bar{\rho}(T)R = (I_{\bar{\rho}} \times T) \circ R$. For Eqn. (35), we have

$$(36) \quad (I_{\rho} \times \bar{T}^*) \circ \bar{R} = \rho(\bar{T}^*)\bar{R} = \rho(R^*)\bar{\rho}\bar{\rho}(T^*\bar{R}')\bar{R} = \rho(R^*)\bar{R}T^*\bar{R}',$$

where we used the definition of \bar{T} for the second equality, and $\bar{R} \in \text{Hom}(\iota, \rho \otimes \bar{\rho})$ for the third equality. But by the conjugate equations $\rho(R^*)\bar{R} = (I_{\rho} \times R^*) \circ (\bar{R} \times I_{\rho}) = I_{\rho}$, and hence $(I_{\rho} \times \bar{T}^*) \circ \bar{R} = T^*\bar{R}' = (T^* \times I_{\bar{\rho}'}) \circ \bar{R}'$. ■

PROPOSITION 263. *The algebra \mathfrak{F}_0 becomes a $*$ -algebra if we define*

$$(37) \quad (A, \rho, \psi)^* := (R^*\bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi),$$

where $\psi \in E(\rho)$, and $(\bar{\rho}, R, \bar{R})$ is a conjugate to ρ .

Proof. We first show that the definition of $*$ is independent of the choice of conjugate to ρ . For this, let $(\bar{\rho}_1, R_1, \bar{R}_1)$ be any other choice. Define $W \in \text{Hom}(\bar{\rho}, \bar{\rho}_1)$ by

$$(38) \quad W := (R^* \times I_{\bar{\rho}_1}) \circ (I_{\bar{\rho}} \times \bar{R}_1) = R^*\bar{\rho}(\bar{R}_1)$$

we have by the conjugate equations

$$W^{-1} := (R_1^* \times I_{\bar{\rho}}) \circ (I_{\bar{\rho}_1} \times \bar{R}) = R_1^*\bar{\rho}_1(\bar{R}).$$

Moreover,

$$\begin{aligned} (R_1^*\bar{\rho}_1(A)^*, \bar{\rho}_1, \mathcal{J}E(\bar{R}_1^*)\psi) &= (R^*W^{-1}\bar{\rho}_1(A)^*, \bar{\rho}_1, \mathcal{J}E(\bar{R}^*(I_{\rho} \times W^*))\psi) \\ &= (R^*\bar{\rho}(A)^*, \bar{\rho}, E(W^{-1})\mathcal{J}E(\bar{R}^*(I_{\rho} \times W^*))\psi) \\ &= (R^*\bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi), \end{aligned}$$

where we used Lemma 260.3 for the final equality.

To see that the definition of $*$ is independent of equivalence classes, suppose that $T \in \text{Hom}(\rho, \rho')$ and $\psi \in E(\rho)$. Then

$$\begin{aligned}
 (AT, \rho, \psi)^* &= (R^*\bar{\rho}(T^*A^*), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\
 &= (R^*\bar{T}^*\bar{\rho}(A^*), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\
 &= (R'^*\bar{\rho}'(A^*)\bar{T}^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\
 &= (R'^*\bar{\rho}'(A^*), \bar{\rho}', E(\bar{T}^*)\mathcal{J}E(\bar{R}^*)\psi) \\
 &= (R'^*\bar{\rho}'(A^*), \bar{\rho}', \mathcal{J}E(\bar{R}^* \circ (I_\rho \times \bar{T}))\psi) \\
 &= (R'^*\bar{\rho}'(A^*), \bar{\rho}', \mathcal{J}E(\bar{R}'^* \circ (T \times I_{\bar{\rho}}))\psi) \\
 &= (R'^*\bar{\rho}'(A^*), \bar{\rho}', \mathcal{J}E(\bar{R}'^*)E(T)\psi) \\
 &= (A, \rho', E(T)\psi)^*,
 \end{aligned}$$

where we used Eqn. (34) for the second equality, the fact that $\bar{T}^* \in \text{Hom}(\bar{\rho}, \bar{\rho}')$ for the third equality, Lemma 260.2 for the fifth equality, and Eqn. (35) for the sixth equality.

We verify that $*$ is involutive:

$$\begin{aligned}
 (A, \rho, \psi)^{**} &= (R^*\bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi)^* \\
 &= (\bar{R}^*\rho(\bar{\rho}(A)R), \rho, \mathcal{J}E(R^*)\mathcal{J}E(\bar{R}^*)\psi) \\
 &= (A\bar{R}^*\rho(R), \rho, \mathcal{J}E(R^*)\mathcal{J}E(\bar{R}^*)\psi) \\
 &= (A, \rho, E((R^* \times I_{\bar{\rho}})(I_{\bar{\rho}} \times \bar{R}))\psi) \\
 &= (A, \rho, \psi),
 \end{aligned}$$

where we used Lemma 260.3 for the penultimate equality, and the conjugate equations for the final equality.

To verify that $*$ is antilinear, let $W_i \in \text{Hom}(\rho_i, \rho)$. Then,

$$\begin{aligned}
 (39) \quad &[(A_1, \rho_1, \psi_1) + (A_2, \rho_2, \psi_2)]^* = (A_1W_1^* + A_2W_2^*, \rho, E(W_1)\psi_1 + E(W_2)\psi_2)^* \\
 &= (R^*\bar{\rho}(W_1A_1^* + W_2A_2^*), \bar{\rho}, \mathcal{J}E(\bar{R}^*)(E(W_1)\psi_1 + E(W_2)\psi_2)).
 \end{aligned}$$

But we may take $R = (\bar{W}_1 \times W_1) \circ R_1 + (\bar{W}_2 \times W_2) \circ R_2$, $\bar{R} = (W_1 \times \bar{W}_1) \circ \bar{R}_1 + (W_2 \times \bar{W}_2) \circ \bar{R}_2$, where $\bar{W}_i \in \text{Hom}(\bar{\rho}_i, \bar{\rho})$ are isometries, $\bar{W}_1\bar{W}_1^* + \bar{W}_2\bar{W}_2^* = I_{\bar{\rho}}$. Then Eqn. (39) becomes

$$\begin{aligned}
 &[(A_1, \rho_1, \psi_1) + (A_2, \rho_2, \psi_2)]^* \\
 &= (R_1^*\bar{\rho}_1(A_1^*)\bar{W}_1^* + R_2^*\bar{\rho}_2(A_2^*)\bar{W}_2^*, \bar{\rho}, \mathcal{J}E(\bar{R}_1^*(I_{\rho_1} \times \bar{W}_1^*))\psi_1 + \\
 &\quad \mathcal{J}E(\bar{R}_2^*(I_{\rho_2} \times \bar{W}_2^*))\psi_2) \\
 &= (R_1^*\bar{\rho}_1(A_1^*)\bar{W}_1^* + R_2^*\bar{\rho}_2(A_2^*)\bar{W}_2^*, \bar{\rho}, E(\bar{W}_1)\mathcal{J}E(\bar{R}_1^*)\psi_1 + E(\bar{W}_2)\mathcal{J}E(\bar{R}_2^*)\psi_2) \\
 &= (A_1, \rho_1, \psi_1)^* + (A_2, \rho_2, \psi_2)^*,
 \end{aligned}$$

using Lemma 260 for the second equality.

Finally, we show that $[(A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2)]^* = (A_2, \rho_2, \psi_2)^*(A_1, \rho_2, \psi_1)^*$. If $\rho = \rho_1 \otimes \rho_2$ and $\rho' = \rho'_1 \otimes \rho'_2$ then we may take $R = (I_{\bar{\rho}_2} \times R_1 \times I_{\rho_2}) \circ R_2$ and $\bar{R} = (I_{\rho_1} \times \bar{R}_2 \times I_{\bar{\rho}_1}) \circ \bar{R}_1$. Thus,

$$\begin{aligned} [(A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2)]^* &= (A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2)^* \\ &= (R_2^* \bar{\rho}_2(R_1^*) \bar{\rho}_2 \bar{\rho}_1(\rho_1(A_2^*) A_1^*), \bar{\rho}_2 \otimes \bar{\rho}_1, \mathcal{J}E(\bar{R}_1^* \circ (I_{\rho_1} \times \bar{R}_2^* \times I_{\bar{\rho}_1}))) \psi_1 \otimes \psi_2) \\ &= (R_2^* \bar{\rho}_2(A_2^*) \bar{\rho}_2(R_1^* \bar{\rho}_1(A_1^*)), \bar{\rho}_2 \otimes \bar{\rho}_1, \mathcal{J}E(\bar{R}_1^* \circ (I_{\rho_1} \times \bar{R}_2^* \times I_{\bar{\rho}_1}))) \psi_1 \otimes \psi_2) \\ &= (R_2^* \bar{\rho}_2(A_2^*) \bar{\rho}_2(R_1^* \bar{\rho}_1(A_1^*)), \bar{\rho}_2 \otimes \bar{\rho}_1, \mathcal{J}E(\bar{R}_2^*) \psi_2 \otimes \mathcal{J}E(\bar{R}_1^*) \psi_1) \\ &= (R_2^* \bar{\rho}_2(A_2^*), \bar{\rho}_2, \mathcal{J}E(\bar{R}_2^*) \psi_2)(R_1^* \bar{\rho}_1(A_1^*), \bar{\rho}_1, \mathcal{J}E(\bar{R}_1^*) \psi_1) \\ &= (A_2, \rho_2, \psi_2)^*(A_1, \rho_1, \psi_1)^*, \end{aligned}$$

where the third equality follows from the fact that $R_1^* \in \text{Hom}(\bar{\rho}_1 \otimes \rho_1, \iota)$, and the fourth equality follows by Lemma 260.4. ■

PROPOSITION 264. *Let $E : \Delta_f \rightarrow \mathcal{SH}_f$ be the embedding functor from the DHR category Δ_f into the strictified category \mathcal{SH}_f of finite dimensional super Hilbert spaces. Then the formula*

$$(40) \quad \alpha_g(A, \rho, \psi) = (A, \rho, g_\rho \psi), \quad A \in \mathfrak{A}, \psi \in E(\rho).$$

defines a group isomorphism $g \mapsto \alpha_g$ from the intrinsic group G of E into $\text{Aut}_{\mathfrak{A}} \mathfrak{F}_0$, the group of $*$ -automorphisms of \mathfrak{F}_0 leaving \mathfrak{A} pointwise fixed.

Proof. Since g is a natural monoidal transformation, $g_\iota = \text{id}_{E(\iota)} = \text{id}_{\mathbb{C}}$. For any $g \in G$, α_g is well defined on \mathfrak{F}_0 since for $S \in \text{Hom}(\rho, \rho')$,

$$\begin{aligned} \alpha_g(AS, \rho, \psi) &= (AS, \rho, g_\rho \psi) = (A, \rho', E(S)g_\rho \psi) \\ &= (A, \rho', g(\rho')E(S)\psi) = \alpha_g(A, \rho', E(S)\psi). \end{aligned}$$

Since $g_\iota = \text{id}_{\mathbb{C}}$, α_g leaves $\mathfrak{A} \subset \mathfrak{F}_0$ pointwise fixed. Each g_ρ is linear so α_g is linear.

$$(A_1, \rho_1, g_{\rho_1} \psi_1)(A_2, \rho_2, g_{\rho_2} \psi_2) = \left(A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, (g_{\rho_1} \otimes g_{\rho_2})(\psi_1 \otimes \psi_2) \right),$$

but $g_{\rho_1 \otimes \rho_2} = g_{\rho_1} \otimes g_{\rho_2}$ so

$$(A_1, \rho_1, g_{\rho_1} \psi_1)(A_2, \rho_2, g_{\rho_2} \psi_2) = \left(A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, g_{\rho_1 \otimes \rho_2}(\psi_1 \otimes \psi_2) \right).$$

Thus,

$$(41) \quad \alpha_g(F_1)\alpha_g(F_2) = \alpha_g(F_1 F_2).$$

To show that α_g is a $*$ -homomorphism, recall that

$$(42) \quad (A, \rho, g_\rho \psi)^* = (R^* \bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)g_\rho \psi).$$

If $\bar{\psi} \in E(\bar{\rho})$ then $E(\bar{R}^*)(g_\rho \psi \otimes g_{\bar{\rho}} \bar{\psi}) = E(\bar{R}^*)((g_\rho \otimes g_{\bar{\rho}})(\psi \otimes \bar{\psi}))$. Furthermore,

$$E(\bar{R}^*)(g_\rho \psi \otimes g_{\bar{\rho}} \bar{\psi}) = g_\iota E(\bar{R}^*)(\psi \otimes \bar{\psi}) = E(\bar{R}^*)(\psi \otimes \bar{\psi}).$$

Hence $g_{\bar{\rho}}^* \mathcal{J}E(\bar{R}^*)g_{\rho} = \mathcal{J}E(\bar{R}^*)$ and since $g_{\bar{\rho}}$ is unitary we get from (42),

$$(A, \rho, g_{\rho}\psi)^* = (R^*\bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)g_{\rho}\psi) = (R^*\bar{\rho}(A)^*, \bar{\rho}, g_{\bar{\rho}}\mathcal{J}E(\bar{R}^*)\psi),$$

so

$$(43) \quad \alpha_g(F^*) = \alpha_g(F)^*, \quad F \in \mathfrak{F}.$$

Equations (41), (43) show that α_g is a $*$ -homomorphism, its inverse is clearly $\alpha_{g^{-1}}$ so α_g defined by Eqn. (40) is an element of $\text{Aut}_{\mathfrak{A}}\mathfrak{F}$. The mapping $g \mapsto \alpha_g$ is clearly a homomorphism.

Since G is a compact group, for every $g \neq e$, there exists a $(H, \pi) \in \text{Rep}_f G$ such that $\pi(g) \neq \text{id}_H$. Since the functor E is an equivalence, in particular essentially surjective, there exists a $\rho \in \text{Obj}(\Delta_f)$ such that $E(\rho)$ is isomorphic to (H, π) . Thus there exists $\psi \in E(\rho)$ such that

$$\pi(g)\psi = g_{\rho}\psi \neq \psi.$$

Defining $F = (I, \rho, \psi)$, we have $\alpha_g(F) \neq F$. This proves injectivity of $g \mapsto \alpha_g$.

It remains to show that $G \mapsto \text{Aut}_{\mathfrak{A}}\mathfrak{F}_0$ is onto. Let $\alpha \in \text{Aut}_{\mathfrak{A}}\mathfrak{F}_0$, $A \in \mathfrak{A}$ and $\psi \in E(\rho) \subset \mathfrak{F}_0$. Let $\Psi = (I, \rho, \psi)$. Then

$$(\alpha(\Psi))A = \alpha(\Psi A) = \alpha(\rho(A)\Psi) = \rho(A)\alpha(\Psi).$$

It is easily checked that this implies that $\alpha(\Psi)$ is of the form (I, ρ, ψ') with $\psi' \in E(\rho)$. Thus $\psi \mapsto \psi'$ is a linear map of $E(\rho)$ into $E(\rho)$ which we denote by g_{ρ} , and it remains to show that $g = (g_{\rho})_{\rho \in \Delta_f}$ is monoidal natural transformation of E . For $S \in \text{Hom}(\rho, \rho')$, we have

$$(S, \rho, g_{\rho}\psi) = \alpha(S, \rho, \psi) = \alpha(I, \rho', E(S)\psi) = (I, \rho', g_{\rho'}E(S)\psi).$$

Hence

$$E(S)g_{\rho}\psi = g_{\rho'}E(S)\psi, \quad \psi \in E(\rho).$$

That is,

$$E(S)g_{\rho} = g_{\rho'}E(S),$$

and $g \in \text{Nat}E$. To check monoidality, choose arbitrary $\psi_i \in E(\rho_i)$ and let $\Psi_i = (I, \rho_i, \psi_i)$. Then,

$$g_{\rho_1 \otimes \rho_2}(\psi_1 \otimes \psi_2) = \alpha(\Psi_1 \Psi_2) = \alpha(\Psi_1)\alpha(\Psi_2) = (g_{\rho_1} \otimes g_{\rho_2})(\psi_1 \psi_2).$$

Thus, $g \in \text{Nat}_{\otimes}E$. It remains to show that g is unitary. For $\psi, \psi' \in E(\rho)$ and $\Psi = (I, \rho, \psi)$, $\Psi' = (I, \rho, \psi')$ we have

$$\langle g_{\rho}\psi, g_{\rho}\psi' \rangle_{E(\rho)} I = \alpha(\Psi)^* \alpha(\Psi') = \alpha(\Psi^* \Psi') = \langle \psi, \psi' \rangle_{E(\rho)} I,$$

where the first and last equalities follow from Prop. 270. Hence g_{ρ} is unitary for each $\rho \in \text{Obj}(\Delta_f)$. Therefore every $\alpha \in \text{Aut}_{\mathfrak{A}}\mathfrak{F}_0$ is of the form α_g with $g \in G = \text{Nat}_{\otimes}E$. ■

DEFINITION 265. Given a double cone O , we define $\mathfrak{F}_0(O)$ to consist of those elements F in \mathfrak{F}_0 such that there exists $A \in \mathfrak{A}(O)$, $\rho \in \text{Obj}(\Delta_f)$ localized in O , and $\psi \in E(\rho)$ with $F = (A, \rho, \psi)$.

PROPOSITION 266. $\mathfrak{F}_0(O)$ is a $*$ -subalgebra of \mathfrak{F}_0 .

Proof. Let $F_1 = (A_1, \rho_1, \psi_1)$ and $F_2 = (A_2, \rho_2, \psi_2)$ be in $\mathfrak{F}_0(O)$. Thus, the A_i can be chosen from $\mathfrak{A}(O)$ and the ρ_i can be chosen localized in O . Since $\rho_1(\mathfrak{A}(O)) \subseteq \mathfrak{A}(O)$, it follows that

$$F_1 F_2 = (A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2),$$

is also in $\mathfrak{F}_0(O)$. By transportability, $\bar{\rho}$ can be chosen localized in O , and in this case $\bar{\rho} \otimes \rho$ is localized in O . By Lemma 155, $R \in \mathfrak{A}(O)$. Hence,

$$F^* = (R^* \bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi),$$

is in $\mathfrak{F}_0(O)$. Similarly, $\mathfrak{F}_0(O)$ is closed under the addition defined in Prop. 257 since ρ can also be chosen localized in O , and then the isometries W_1, W_2 are in $\mathfrak{A}(O)$ (by Lemma 155). ■

PROPOSITION 267. The action of G on \mathfrak{F}_0 leaves $\mathfrak{F}_0(O)$ globally fixed.

Proof. If $F \in \mathfrak{F}_0(O)$ then $F = (A, \rho, \psi)$ for some $A \in \mathfrak{A}(O)$ and ρ localized in O . Then clearly $\alpha_g(F) = (A, \rho, g_\rho \psi)$ is in $\mathfrak{F}_0(O)$. ■

NOTE 268. Having defined an action of the supergroup (G, k) , the element $k \in G$ induces a \mathbb{Z}_2 grading on \mathfrak{F}_0 and on the local algebras $\mathfrak{F}_0(O)$.

PROPOSITION 269. The field net \mathfrak{F}_0 satisfies normal commutation relations. That is, if O_1 and O_2 are spacelike, and $F_i \in \mathfrak{F}(O_i)$ are such that

$$\alpha_k(F_i) = \sigma_i F_i,$$

then

$$F_1 F_2 = (-1)^{(1-\sigma_1)(1-\sigma_2)/4} F_2 F_1.$$

Proof. Choose $F_i = (A_i, \rho_i, \psi_i)$ with $A_i \in \mathfrak{A}(O_i)$ and ρ_i localized in O_i . Then $A_1 A_2 = A_2 A_1$, $\rho_1(A_2) = A_2$, $\rho_2(A_1) = A_1$, and $\varepsilon_{\rho_1, \rho_2} = \text{id}_{\rho_1 \otimes \rho_2}$. In view of the way G acts on \mathfrak{F}_0 we have

$$\sigma_i(A_i, \rho_i, \psi_i) = \alpha_k(A_i, \rho_i, \psi_i) = (A_i, \rho_i, k_{\rho_i} \psi_i),$$

and hence $k_{\rho_i} \psi_i = \sigma_i \psi_i$. That is, ψ_i is homogeneous and $\omega(\psi_i) = \sigma_i$. Furthermore, since E is a symmetric functor $E(\varepsilon_{\rho_1, \rho_2}) = \Sigma_{E(\rho_1), E(\rho_2)}$, where $\Sigma_{H, H'}$ is the symmetry on \mathcal{SH}_f and therefore

$$\Sigma_{H, H'}(\psi_1 \otimes \psi_2) = (-1)^{(1-\sigma_1)(1-\sigma_2)/4} (\psi_2 \otimes \psi_1).$$

Hence

$$\begin{aligned}
 F_1 F_2 &= (A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2) \\
 &= (A_1 A_2 \varepsilon_{\rho_2, \rho_1}, \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2) \\
 &= (A_1 A_2, \rho_2 \otimes \rho_1, E(\varepsilon_{\rho_2, \rho_1})(\psi_1 \otimes \psi_2)) \\
 &= (A_2 \rho_2(A_1), \rho_2 \otimes \rho_1, E(\varepsilon_{\rho_2, \rho_1})(\psi_1 \otimes \psi_2)) \\
 &= (A_2 \rho_2(A_1), \rho_2 \otimes \rho_1, \Sigma_{E(\rho_2), E(\rho_1)}(\psi_1 \otimes \psi_2)) \\
 &= (-1)^{(1-\sigma_1)(1-\sigma_2)/4} (A_2 \rho_2(A_1), \rho_2 \otimes \rho_1, \psi_2 \otimes \psi_1) \\
 &= (-1)^{(1-\sigma_1)(1-\sigma_2)/4} F_2 F_1.
 \end{aligned}$$

■

PROPOSITION 270. For all $\Psi = (I, \rho, \psi)$, $\Psi' = (I, \rho, \psi')$ with $\psi, \psi' \in E(\rho)$ we have

$$(44) \quad \Psi A = \rho(A) \Psi,$$

$$(45) \quad \Psi^* \Psi' = \langle \psi, \psi' \rangle_{E(\rho)}.$$

For any orthonormal basis $\{\psi_i : i = 1, \dots, n\}$ of $E(\rho)$, we have

$$(46) \quad \sum_{i=1}^n \Psi_i \Psi_i^* = I.$$

Proof.

$$(I, \rho, \psi)(A, \iota, 1) = (\rho(A), \rho, \psi) = (\rho(A), \iota, 1)(I, \rho, \psi),$$

whence (44). For (45), we check:

$$\begin{aligned}
 (I, \rho, \psi)^*(I, \rho, \psi') &= (R^*, \bar{\rho} \otimes \rho, (\mathcal{J}E(\bar{R}^*)\psi) \otimes \psi') \\
 &= (I, \iota, E(R^*)((\mathcal{J}E(\bar{R}^*)\psi) \otimes \psi')).
 \end{aligned}$$

Since $\mathcal{J}E(\bar{R}^*) : E(\rho) \rightarrow E(\bar{\rho})$ and $E(R^*) : E(\bar{\rho}) \otimes E(\rho) \rightarrow \mathbb{C}$, it follows that $E(R^*)((\mathcal{J}E(\bar{R}^*)\psi) \otimes \psi')$ is a complex number. In fact, by the definition of \mathcal{J} and Lemma 260.3,

$$\begin{aligned}
 E(R^*)((\mathcal{J}E(\bar{R}^*)\psi) \otimes \psi') &= \langle \mathcal{J}E(R^*) \circ \mathcal{J}E(\bar{R}^*)\psi, \psi' \rangle_{E(\rho)} \\
 &= \langle \mathcal{J}E((R^* \times I_\rho) \circ (I_\rho \times \bar{R}^*))\psi, \psi' \rangle_{E(\rho)} \\
 &= \langle \psi, \psi' \rangle_{E(\rho)},
 \end{aligned}$$

where the final equality follows by the conjugate equations. So, combining the previous two equations we have

$$(I, \rho, \psi)^*(I, \rho, \psi') = (I, \iota, \langle \psi, \psi' \rangle_{E(\rho)}) = \langle \psi, \psi' \rangle_{E(\rho)} (I, \iota, 1).$$

For Eqn. (46), we have

$$\begin{aligned} \sum_i (I, \rho, \psi_i)(I, \rho, \psi_i)^* &= \left(\rho(R)^*, \rho \otimes \bar{\rho}, \sum \psi_i \otimes \mathcal{J}E(\bar{R}^*)\psi_i \right) \\ &= (\rho(R)^*, \rho \otimes \bar{\rho}, E(\bar{R})1) = (\rho(R)^*\bar{R}, \iota, 1) = (I, \iota, 1), \end{aligned}$$

where the second equality follows from the definition of \mathcal{J} and the final equality follows by the conjugate equations. ■

10.3 Completion of the field net

We now construct a representation (\mathcal{H}, π) of the $*$ -algebra \mathfrak{F}_0 , and show that $\pi|_{\mathfrak{A}}$ has a nontrivial subrepresentation equivalent to the GNS representation induced by the vacuum state ω_0 . We do so by extending the state ω_0 from \mathfrak{A} to \mathfrak{F}_0 , and then by taking the GNS representation. In order to extend the state ω_0 from \mathfrak{A} to \mathfrak{F}_0 , it suffices to show that there is a positive linear map $m : \mathfrak{F}_0 \rightarrow \mathfrak{A}$.

NOTE 271. Let $\rho \in \text{Obj}(\Delta_f)$. Since Δ_f is semisimple (see Prop. 218), ρ is a finite direct sum $\rho = \rho_1 \oplus \dots \oplus \rho_n$ of irreducible objects in $\text{Obj}(\Delta_f)$. Therefore, there is a projection $P_\iota^\rho \in \text{End}(\rho)$ onto the direct sum of those irreducibles in this decomposition that are isomorphic to ι .

PROPOSITION 272. *Given $(A, \rho, \psi) \in \mathfrak{F}_0$, define*

$$(47) \quad m(A, \rho, \psi) := (AP_\iota^\rho, \rho, \psi).$$

Then $m : \mathfrak{F}_0 \rightarrow \mathfrak{A}$ is a faithful positive linear projection from \mathfrak{F}_0 onto \mathfrak{A} . Further,

$$(48) \quad m(AF) = Am(F), \quad A \in \mathfrak{A}, F \in \mathfrak{F}_0.$$

Proof. We first show that m is well defined. If $T \in \text{Hom}(\rho, \rho')$ then $TP_\iota^\rho = P_\iota^{\rho'}TP_\iota^\rho = P_\iota^{\rho'}T$, hence

$$m(AT, \rho, \psi) = (ATP_\iota^\rho, \rho, \psi) = (AP_\iota^{\rho'}T, \rho, \psi) = (AP_\iota^{\rho'}, \rho', E(T)\psi) = m(A, \rho', E(T)\psi),$$

as required. m is clearly linear and satisfies Eqn. (48). We now show that m is positive. First, since ρ has finite dimension, ρ contains at most finitely many copies of the vacuum representation. Thus, $P_\iota^\rho = \sum_i S_i S_i^*$ where $S_i \in \text{Hom}(\iota, \rho)$ and $S_i^* S_j = \delta_{ij} \text{id}_\iota$. Thus,

$$m(A, \rho, \psi) = (AP_\iota^\rho, \rho, \psi) = \sum_i (AS_i, \iota, E(S_i^*)\psi).$$

However, $E(S_i^*)\psi = \lambda_i 1$ so that

$$m(A, \rho, \psi) = \sum_i \lambda_i (AS_i, \iota, 1) \in \mathfrak{A}.$$

Since each $\rho \in \text{Obj}(\Delta_f)$ is a finite direct sum of irreducible objects (Prop. 218), any $F \in \mathfrak{F}_0$ may be written as a finite sum $F = \sum_i F_i$, $F_i = (A_i, \rho_i, \psi_i)$, where $\psi_i \in E(\rho_i)$ with ρ_i irreducible and pairwise inequivalent. Thus,

$$m(F^*F) = \sum_{i,j} m(F_i^*F_j) = \sum_i m(F_i^*F_i).$$

Hence, to show that m is positive and faithful, it suffices to consider $m(F^*F)$ with $F = (A, \rho, \psi)$, $\psi \in E(\rho)$ and ρ irreducible. In this case,

$$(A, \rho, \psi)^*(A, \rho, \psi) = (R^*\bar{\rho}(A^*A), \bar{\rho} \otimes \rho, \mathcal{J}E(\bar{R}^*)(\psi \otimes \psi)).$$

Using $P_t^{\bar{\rho} \otimes \rho} = \|RR^*\|^{-1}RR^* = d(\rho)^{-1}RR^*$, we have

$$\begin{aligned} d(\rho)m(F^*F) &= (R^*\bar{\rho}(A^*A)RR^*, \bar{\rho} \otimes \rho, \mathcal{J}E(\bar{R}^*)\psi \otimes \psi) \\ &= (R^*\bar{\rho}(A^*A)R, \iota, E(R^*)\mathcal{J}E(\bar{R}^*)\psi \otimes \psi). \end{aligned}$$

Now,

$$E(R^*)\mathcal{J}E(\bar{R}^*)(\psi \otimes \psi) = \langle \mathcal{J}E(R^*)\mathcal{J}E(\bar{R}^*)\psi, \psi \rangle_{E(\rho)},$$

hence by Lemma 260,

$$\begin{aligned} d(\rho)m(F^*F) &= R^*\bar{\rho}(A^*A)R \langle E((\bar{R}^* \times I_\rho) \circ (I_\rho \times R))\psi, \psi \rangle_{E(\rho)} \\ &= R^*\bar{\rho}(A^*A)R \langle \psi, \psi \rangle_{E(\rho)}. \end{aligned}$$

Thus, $m(F^*F) \geq 0$ and $m(F^*F) = 0$ implies $\psi = 0$ or $\bar{\rho}(A)R = 0$. But $\bar{\rho}(A)R = 0$ only if

$$0 = \bar{R}^* \rho \bar{\rho}(A) \rho(R) = A \bar{R}^* \rho(R) = A.$$

Thus $m(F^*F) = 0$ implies $F = 0$, and m is a faithful positive linear projection from \mathfrak{F}_0 onto \mathfrak{A} . ■

LEMMA 273. *Let P_0^ρ be the projection in $\text{End}(E(\rho))$ onto the subspace of G invariant vectors with respect to the action $\pi_\rho(g) = g_\rho$. Then $E(P_t^\rho) = P_0^\rho$. Furthermore, the conditional expectation m is G -invariant, i.e. $m(\alpha_g(F)) = m(F)$ for all $g \in G$ and $F \in \mathfrak{F}_0$.*

Proof. Recall that if (H, π) is an irreducible representation of a compact group G and π is not the trivial representation, then H contains no G -invariant vectors. If $\rho = \bigoplus \rho_i$ with ρ_i irreducible, then the previous observation implies that the G -invariant vectors in $E(\rho)$ are precisely those in the image of $E(P_t^\rho)$. Thus $E(P_t^\rho) = P_0^\rho$, implying $m(F) = \alpha_g(m(F))$. Furthermore,

$$\begin{aligned} m\alpha_g(A, \rho, \psi) &= m(A, \rho, g_\rho\psi) = (AP_t^\rho, \rho, g_\rho\psi) = (A, \rho, P_0^\rho g_\rho\psi) \\ &= (A, \rho, g_\rho P_0^\rho\psi) = (A, \rho, P_0^\rho\psi) \\ &= (AP_t^\rho, \rho, \psi) = m(A, \rho, \psi). \end{aligned}$$

■

In view of Prop. 272, $\omega_0 \circ m$ is a faithful state on the $*$ -algebra \mathfrak{F}_0 . Let (\mathcal{H}, π) be the GNS representation of \mathfrak{F} induced by $\omega_0 \circ m$, let \mathfrak{F} be the norm closure of $\pi(\mathfrak{F}_0)$, and let $\mathfrak{F}(O)$ be the *weak* closure of $\pi(\mathfrak{F}_0(O))$. It is clear that \mathfrak{F} is the C^* -inductive limit of the net $O \mapsto \mathfrak{F}(O)$. Since $\omega_0 \circ m$ is G -invariant by Lemma 273, there is a unitary representation U of G on \mathcal{H} implementing the automorphisms α_g of \mathfrak{F}_0 :

$$\pi(\alpha_g(F)) = U(g)\pi(F)U(g)^*, \quad g \in G, F \in \mathfrak{F}_0,$$

and therefore it extends to \mathfrak{F} . Since $g \mapsto \alpha_g$ is injective, U is injective.

DEFINITION 274. Let $\sigma \in \hat{G}$ be an irreducible character of G . Define a map \mathcal{E}_σ on $\mathfrak{B}(\mathcal{H})$ by

$$\mathcal{E}_\sigma(A) = \int_G \overline{\sigma(g)}U(g)AU(g)^* d\mu(g),$$

where μ is the Haar measure on $U(G)$.

NOTE 275. Let $F = (A, \rho, \psi) \in \mathfrak{F}_0$. Since the $U(g)$ implements α_g we have

$$\begin{aligned} \mathcal{E}_\sigma(\pi(F)) &= \int_G \overline{\sigma(g)}\left(\pi(\alpha_g(A, \rho, \psi))\right)d\mu(g) = \int_G \overline{\sigma(g)}\pi(A, \rho, g_\rho\psi)d\mu(g) \\ (49) \qquad &= \pi(A, \rho, P_\sigma^\rho\psi), \end{aligned}$$

where $P_\sigma^\rho \in \text{End}(E(\rho))$ is the orthogonal projection onto the subspace transforming according to the irreducible representation σ . Since G is compact, \mathcal{E}_σ is strongly continuous. Note that $\mathcal{E}_0(\pi(F)) = \pi[m(A, \rho, \psi)]$.

LEMMA 276. $\mathfrak{F}_0(O)\mathfrak{A} = \mathfrak{F}_0$.

Proof. Let $(A, \rho, \psi) \in \mathfrak{F}_0$. Since ρ is transportable, there is a unitary $T \in \text{Hom}(\rho, \rho')$ with ρ' localized in O . Then

$$(A, \rho, \psi) = (AT^*, \rho', E(T)\psi) = (AT^*, \iota, 1)(I, \rho', E(T)\psi) = BF,$$

where $B \in \mathfrak{A}$ and $F \in \mathfrak{F}_0(O)$. Hence $\mathfrak{A}\mathfrak{F}_0(O) = \mathfrak{F}_0$. Since $\mathfrak{A}, \mathfrak{F}_0(O)$ and \mathfrak{F}_0 are $*$ -algebras, $\mathfrak{F}_0(O)\mathfrak{A} = \mathfrak{F}_0$. ■

THEOREM 277. $(\pi, \mathcal{H}, \mathfrak{F}, (G, k))$ is a field system with gauge symmetry for (\mathfrak{A}, ω_0) with normal commutation relations (in the sense of Definitions 229 and 245).

Proof. It is obvious that $\mathfrak{F}(O)$ is a G -stable von Neumann subalgebra of \mathfrak{F} . Also the net $O \mapsto \mathfrak{F}(O)$ satisfies normal commutation relations. We now run through the individual conditions in Definition 229.

(γ) We need to show that the fixed point algebra of $\mathfrak{F}(O)$ under the G action is $\pi(\mathfrak{A}(O))$. First note that $\mathcal{E}(\pi(\mathfrak{F}_0(O))) = \pi(m(\mathfrak{F}_0(O)))$. Thus,

$$\mathfrak{F}(O)^G = \mathcal{E}(\mathfrak{F}(O)) = \mathcal{E}\left(\overline{\pi(\mathfrak{F}_0(O))}\right) = \overline{\mathcal{E}(\pi(\mathfrak{F}_0(O)))} = \overline{\pi(m(\mathfrak{F}_0(O)))} = \overline{\pi(\mathfrak{A}(O))}.$$

The third equality follows by the normality of \mathcal{E} , and the last equality is due to the fact that m is a conditional expectation from \mathfrak{F}_0 to \mathfrak{A} .

(δ) Let $j : \mathfrak{F}_0 \rightarrow \mathcal{H}$ be the inclusion mapping derived from the GNS representation of $\omega_0 \circ m$. Since $\overline{j(\mathfrak{A})} = \mathcal{H}_0$ we have

$$\overline{\mathfrak{F}(O)\mathcal{H}_0} = \overline{\pi(\mathfrak{F}_0(O))\mathcal{H}_0} = \overline{\pi(\mathfrak{F}_0(O))j(\mathfrak{A})} = \overline{j(\mathfrak{F}_0(O)\mathfrak{A})} = \overline{j(\mathfrak{F}_0)} = \mathcal{H}.$$

(ε) Let O_1 and O_2 be spacelike separated. The subalgebra $\mathfrak{A}(O_1)$ of \mathfrak{F}_0 is pointwise invariant under the gauge transformations. In particular, $\alpha_k(A) = A$ for all $A \in \mathfrak{A}(O)$, i.e. elements of $\mathfrak{A}(O_1)$ are purely Bosonic. Therefore relative locality follows by normality of the commutation relations (Prop. 269).

Now we claim that $\text{Aut}_{\mathfrak{A}}\mathfrak{F} = G$. By Eqn. (49), $\mathcal{E}_\sigma(\pi(\mathfrak{F}_0))$ is isomorphic as a Banach space to $\mathfrak{A} \otimes P_\sigma^0 E(\rho)$, and so is a closed subspace of \mathfrak{F} , and so

$$\mathcal{E}_\sigma(\mathfrak{F}) = \mathcal{E}_\sigma(\overline{\pi(\mathfrak{F}_0)}) = \overline{\mathcal{E}_\sigma(\pi(\mathfrak{F}_0))} = \mathcal{E}_\sigma(\pi(\mathfrak{F}_0)).$$

Since for any $F \in \mathfrak{F}$ we have $F = \sum_{\sigma \in \hat{G}} \mathcal{E}_\sigma(F)$, and $\mathcal{E}_\sigma(F) \in \pi(\mathfrak{F}_0)$, it follows that an element $F \in \mathfrak{F}$ is in $\pi(\mathfrak{F}_0)$ if and only if $\mathcal{E}_\sigma(F) \neq 0$ for only finitely many $\sigma \in \hat{G}$. Together with linearity of α , this implies that $\alpha(\pi(\mathfrak{F}_0)) \subseteq \pi(\mathfrak{F}_0)$. Thus there exists a $g \in G$ such that $\alpha|_{\pi(\mathfrak{F}_0)} = \alpha_g$ (by Prop. 264). Since α_g is continuous and $\pi(\mathfrak{F}_0)$ is dense in \mathfrak{F} , α is the unique extension of α_g to \mathfrak{F} . ■

10.4 Poincaré covariance of the field net

Covariance considerations have played no prominent role in the DHR theory of Section 7 or in the above reconstruction of a field net \mathfrak{F} . We now show that the latter is Poincaré covariant if the underlying DHR sectors are. (Recall from Remark 224 that under favorable circumstances we have $\Delta_{fc} = \Delta_f$.)

THEOREM 278. *If in the construction of the field net \mathfrak{F} we start from the category Δ_{fc} instead of Δ_f , the field net constructed above is covariant under an automorphic action of \hat{P} . This action is implemented by a positive energy representation on the GNS representation space of \mathfrak{F} corresponding to the state $\omega_0 \circ m$.*

Proof. Let $\beta_h = \text{Ad}U(h)$ be the action of \mathcal{P} on \mathfrak{A} . Recall from Note 221 that $\rho_h = \beta_h \circ \rho \circ \beta_h^{-1}$ and $X_\rho(h) \equiv U(h)U_\rho(h)^* \in \text{Hom}(\rho, \rho_h)$ for all $h \in \hat{P}$. We define an action $\hat{\beta}$ of \hat{P} on \mathfrak{F}_0 by

$$\begin{aligned} \hat{\beta}_h((A, \rho, \psi)) &\equiv (\beta_h(A), \rho_h, E(X_\rho(h))\psi) \\ (50) \qquad \qquad &= (\beta_h(A)X_\rho(h), \rho, \psi) = (U(h)AU_\rho(h)^*, \rho, \psi). \end{aligned}$$

Let $\rho, \rho' \in \Delta_{fc}$ and $T \in \text{Hom}(\rho, \rho')$. Then $\beta_h(T) \in \text{Hom}(\rho_h, \rho'_h)$, and $TU_\rho(h) = U_{\rho'}(h)T$, cf. Section 8.4. Thus,

$$\begin{aligned} \beta_h(T)X_\rho(h) &= (U(h)TU(h)^*)(U(h)U_\rho(h)^*) = U(h)TU_\rho(h)^* \\ &= U(h)U_{\rho'}(h)T = X_{\rho'}(h)T, \end{aligned}$$

Using this equation, we compute

$$\begin{aligned} \widehat{\beta}_h((AT, \rho, \psi)) &= (\beta_h(AT), \rho_h, E(X_\rho(h))\psi) = (\beta_h(A), \rho'_h, E(\beta_h(T)X_\rho(h))\psi) \\ &= (\beta_h(A), \rho'_h, E(X_{\rho'}(h)T)\psi) = \widehat{\beta}_h((A, \rho', E(T)\psi)), \end{aligned}$$

thus $\widehat{\beta}_g$ is well defined. Let $i : A \mapsto (A, \iota, \mathbf{1})$ be the inclusion of \mathfrak{A} in \mathfrak{F} . Then $\widehat{\beta}_h \circ i = i \circ \beta_g$, thus $\widehat{\beta}_g$ extends β_g . If $F \in \mathfrak{F}(O)$ then there exists a representation $F = (A, \rho, \psi)$ with $A \in \mathfrak{A}(O)$ and $\rho \in \Delta(O)$. Now it is evident from the definition that $\widehat{\beta}_h(F) \in \mathfrak{F}(hO)$. That $g \mapsto \widehat{\beta}_g$ is a group homomorphism is obvious from the r.h.s. of Eqn. (50). Now,

$$\begin{aligned} \widehat{\beta}_g((A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2)) &= \widehat{\beta}_g((A_1\rho_1(A_2), \rho_1\rho_2, \psi_1 \otimes \psi_2)) \\ &= (U(h)A_1\rho_1(A_2)U_{\rho_1\rho_2}(h)^*, \rho_1\rho_2, \psi_1 \otimes \psi_2) \\ &= (\beta_h(A_1)\rho_{1,h}(\beta_h(A_2))U(h)U_{\rho_1\rho_2}(h)^*, \rho_1\rho_2, \psi_1 \otimes \psi_2) \\ &= (\beta_h(A_1)\rho_{1,h}(\beta_h(A_2))X_{\rho_1\rho_2}(h), \rho_1\rho_2, \psi_1 \otimes \psi_2) \\ &= (\beta_h(A_1)\rho_{1,h}(\beta_h(A_2))X_{\rho_1}(h)\rho_1(X_{\rho_2}(h)), \rho_1\rho_2, \psi_1 \otimes \psi_2) \\ &= (\beta_h(A_1)X_{\rho_1}(h)\rho_1(\beta_h(A_2)X_{\rho_2}(h)), \rho_1\rho_2, \psi_1 \otimes \psi_2) \\ &= (U(h)A_1U_{\rho_1}(h)^*\rho_1(U(h)A_2U_{\rho_2}(h)^*), \rho_1\rho_2, \psi_1 \otimes \psi_2) \\ &= (U(h)A_1U_{\rho_1}(h)^*, \rho_1, \psi_1)(U(h)A_2U_{\rho_2}(h)^*, \rho_2, \psi_2) \\ &= \widehat{\beta}_g((A_1, \rho_1, \psi_1))\widehat{\beta}_g((A_2, \rho_2, \psi_2)), \end{aligned}$$

where the fifth equality is due to Eqn. (27). Thus $\widehat{\beta}_g$ is an algebra homomorphism.

Let $\rho \in \Delta_{fc}$ and choose a conjugate $(\bar{\rho}, R, \bar{R})$. Since the trivial morphism ι is covariant with $X_\iota = \text{id}_L$, applying Eqn. (50) with $T = R^* \in \text{Hom}(\bar{\rho}\rho, \iota)$ we get $R^* = \beta_h(R^*)X_{\bar{\rho}\rho}(h) = \beta_h(R^*)X_{\bar{\rho}}(h)\bar{\rho}(X_\rho(h))$, where we used Eqn. (27) again. This is equivalent to

$$(51) \quad R^*\bar{\rho}(X_\rho(h)^*) = \beta_h(R^*)X_{\bar{\rho}}(h),$$

which will be used below. Now we compute

$$\begin{aligned} (\widehat{\beta}_h(A, \rho, \psi))^* &= (U(h)AU_\rho(h)^*, \rho, \psi)^* \\ &= (R^*\bar{\rho}(U(h)AU_\rho(h)^*)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= (R^*\bar{\rho}(U_\rho(h)A^*U(h)^*), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= (R^*\bar{\rho}(U_\rho(h)U(h)^*\beta_h(A^*)), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= (R^*\bar{\rho}(X_\rho(h)^*\beta_h(A^*)), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= (\beta_h(R^*)X_{\bar{\rho}}(h)\bar{\rho}(\beta_h(A)^*), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= (U(h)R^*U_{\bar{\rho}}(h)^*\bar{\rho}(\beta_h(A)^*), \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= (U(h)R^*\bar{\rho}(A)^*U_{\bar{\rho}}(h)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi) \\ &= \widehat{\beta}_h((R^*\bar{\rho}(A)^*, \bar{\rho}, \mathcal{J}E(\bar{R}^*)\psi)) \\ &= \widehat{\beta}_h((A, \rho, \psi)^*), \end{aligned}$$

thus $\widehat{\beta}_h$ is a $*$ -homomorphism. (In the sixth equality we used Eqn. (51).)

In view of

$$\begin{aligned} \widehat{\beta}_h((A, \rho, \psi)) &= (U(h)AU_\rho(h)^*, \rho, \psi) \\ \alpha_g((A, \rho, \psi)) &= (A, \rho, \pi_{E(\rho)}(g)\psi) \end{aligned}$$

it is clear that $\widehat{\beta}_h \circ \alpha_g = \alpha_g \circ \widehat{\beta}_h$ for all $g \in G, h \in \widehat{\mathcal{P}}$. In view of $\pi \circ m = \mathcal{E}_0 \circ \pi$, we have $\omega_0 \circ \beta_h \circ m = \omega_0 \circ m$. Thus the vacuum state of \mathfrak{F} is $\widehat{\mathcal{P}}$ -invariant, and $\widehat{\mathcal{P}}$ is unitarily implemented in the GNS representation. ■

10.5 Uniqueness of the field net

In the present section we have shown that, given a fiber functor $E : \Delta_f(\mathfrak{A}) \rightarrow \mathcal{SH}$, there exists a field net with normal commutation relations that is complete, i.e. creates all representations in $\Delta_f(\mathfrak{A})$ from the vacuum. We call this the Roberts field net and denote it by \mathfrak{F}_E^R . We first consider the dependence of this construction on the functor E .

PROPOSITION 279. *Let $E_1, E_2 : \Delta_f \rightarrow \mathcal{H}$ be two fiber functors. Then the Roberts field nets $\mathfrak{F}_{E_1}^R, \mathfrak{F}_{E_2}^R$ constructed from them are unitarily equivalent.*

Proof. By Theorem 373 from the appendix, there exists a unitary monoidal natural isomorphism $\alpha : E_1 \rightarrow E_2$. Based on this we define a map $\gamma : \mathfrak{F}_{0,1}^R \rightarrow \mathfrak{F}_{0,2}^R$ by $\gamma : (A, \rho, \psi) \mapsto (A, \rho, \alpha_\rho \psi)$. This makes sense since $\psi \in E_1(\rho)$ and $\alpha \in \text{Hom}(E_1(\rho), E_2(\rho))$. γ is well defined since, for $T \in \text{Hom}(\rho, \rho')$, we have

$$\begin{aligned} \gamma(AT, \rho, \psi) &= (AT, \rho, \alpha_\rho \psi) \\ &= (A, \rho', E_2(T) \circ \alpha_\rho \psi) \\ &= (A, \rho', \alpha_\rho \circ E_1(T)\psi) \\ &= \gamma(A, \rho', E_1(T)\psi). \end{aligned}$$

That γ is an algebra homomorphism follows from

$$\begin{aligned} \gamma((A_1, \rho_1, \psi_1))\gamma((A_2, \rho_2, \psi_2)) &= (A_1, \rho_1, \alpha_{\rho_1} \psi_1)(A_2, \rho_2, \alpha_{\rho_2} \psi_2) \\ &= (A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \alpha_{\rho_1} \psi_1 \otimes \alpha_{\rho_2} \psi_2) \\ &= (A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \alpha_{\rho_1 \otimes \rho_2}(\psi_1 \otimes \psi_2)) \\ &= \gamma((A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, \psi_1 \otimes \psi_2)) \\ &= \gamma((A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2)), \end{aligned}$$

where we have used monoidality $\alpha_{\rho_1 \otimes \rho_2} = \alpha_{\rho_1} \otimes \alpha_{\rho_2}$ of α . Since an inverse can be obtained using the natural isomorphism α^* , γ is an isomorphism between the field algebras $\mathfrak{F}_{0,1}^R$ and $\mathfrak{F}_{0,2}^R$. It clearly respects the local structure, i.e. maps $\mathfrak{F}_{0,1}^R(O)$ to $\mathfrak{F}_{0,2}^R(O)$.

Next we claim that $m_2 \circ \gamma = \gamma \circ m_1$, where m_1, m_2 are the projections defined earlier. Namely,

$$\begin{aligned} m_2 \circ \gamma((A, \rho, \psi)) &= m_2((A, \rho, \alpha_\rho \psi)) = (AP_\ell^\rho, \rho, \alpha_\rho \psi) = \\ &= \gamma((AP_\ell^\rho, \rho, \psi)) = \gamma \circ m_1((A, \rho, \psi)). \end{aligned}$$

This implies that the states $\omega_0 \circ m_1$ and $\omega_0 \circ m_2 \circ \gamma$ on $\mathfrak{F}_{0,1}^R$ coincide, and therefore the isomorphism $\gamma : \mathfrak{F}_{0,1}^R \rightarrow \mathfrak{F}_{0,2}^R$ extends to a unitary equivalence of the norm completions in the GNS representations. ■

In order to study an arbitrary complete normal field net \mathfrak{F} , not a priori of the form $\overline{\mathfrak{F}^R}^{\|\cdot\|}$, we use the following

PROPOSITION 280. *Let \mathfrak{F} be a complete normal field net for the observable net \mathfrak{A} . Then there exists a strict tensor functor $E_{\mathfrak{F}} : \Delta_f(\mathfrak{A}) \rightarrow \mathcal{SH}_f$ to the category of finite dimensional super Hilbert spaces. On the objects, $E_{\mathfrak{F}}$ is given by the vector space*

$$E_{\mathfrak{F}}(\rho) := \{F \in \mathfrak{F} \mid F\pi_0(A) = \pi_0(\rho(A))F \quad \forall A \in \mathfrak{A}\}.$$

The inner product is given by $\langle F, F' \rangle \mathbf{1} = F^*F'$ and the \mathbb{Z}_2 -grading by the action of $k \in G$. For irreducible $\rho, \rho' \in \Delta_f$, we have $E(\varepsilon(\rho, \rho')) = \pm \sum_{i,j} \psi'_i \psi_j \psi_i^* \psi_j^*$, where $\{\psi_i, i = 1, \dots, d(\rho)\}$ and $\{\psi'_i, i = 1, \dots, d(\rho')\}$ are orthonormal bases of $E(\rho)$ and $E(\rho')$, respectively, and the minus sign appears iff ρ and ρ' are both fermionic.

Proof. (In this proof we write E instead of $E_{\mathfrak{F}}$.) For $s \in \text{Hom}(\rho, \rho')$ we define $E(s) = \pi_0(s) \in \mathfrak{F}$. For $F \in E(\rho)$ we have $\pi_0(s)F\pi_0(A) = \pi_0(s)\pi_0(\rho(A))F = \pi_0(s\rho(A))F = \pi_0(\rho'(A))sF = \pi_0(\rho'(A))\pi_0(s)F$ for all $A \in \mathfrak{A}$, thus $\pi_0(s)F \in E(\rho')$ and E is a functor. If $F, F' \in E(\rho)$ then $F^*F' \in \mathfrak{F} \cap \mathfrak{A}' = \mathbb{C}\mathbf{1}$, allowing us to define $\langle F, F' \rangle \mathbf{1} = F^*F'$. Let $s \in \text{Hom}(\rho, \rho')$ and $F \in E(\rho), F' \in E(\rho')$. Then

$$\begin{aligned} \langle F', E(s)F \rangle &= \langle F', \pi_0(s)F \rangle = F'^* \pi_0(s)F = (\pi_0(s)^* F')^* F = \\ &= \langle \pi_0(s^*) F', F \rangle = \langle E(s^*) F', F \rangle, \end{aligned}$$

where we have used that π_0 is $*$ -preserving, shows that E is $*$ -preserving. By Section 9 we have $E(\rho)E(\rho') = E(\rho \otimes \rho')$. If $S_i \in \text{Hom}(\rho_i, \rho'_i), F_i \in E(\rho_i)$ then

$$E(S_1 \times S_2)F_1F_2 = \pi_0(S_1\rho_1(S_2))F_1F_2 = \pi_0(S_1)F_1\pi_0(S_2)F_2 \in E(\rho'_1 \otimes \rho'_2),$$

thus $E(S_1 \times S_2) = E(S_1) \otimes E(S_2)$, thus E is a strict tensor functor. Completeness of the field net together with the discussion in Section 9 implies that E is faithful and satisfies $\dim E(\rho) = d(\rho)$. (The latter follows also by Proposition 344 of the appendix.) Finally, let $F \in E(\rho), F' \in E(\rho')$ be of norm one. Now let $\rho, \rho' \in \Delta_f$, and let $\psi_i, i = 1, \dots, d(\rho)$ and $\psi'_i, i = 1, \dots, d(\rho')$ be orthonormal bases of $E(\rho)$ and $E(\rho')$, respectively. Then

$$\tilde{c}(\rho, \rho') = \sum_{i,j} \psi'_i \psi_j \psi_i^* \psi_j^*$$

is in \mathfrak{F}^G and independent of the chosen bases. Furthermore, $\tilde{c}(\rho, \rho') \in \text{Hom}(\rho \otimes \rho', \rho' \otimes \rho)$. The functoriality of E that was proven above implies that $\tilde{c}(\rho, \rho')$ is natural in both arguments. If now $\{\rho'' \in \Delta_f\}$ and $\psi''_k, k = 1, \dots, d(\rho'')$ is an

orthonormal basis in $E(\rho'')$, then $\{\psi'_j \psi''_k\}$ is an orthonormal basis in $E(\rho' \otimes \rho'')$, thus

$$\begin{aligned} \tilde{c}(\rho, \rho' \otimes \rho'') &= \sum_{i,j,k} \psi'_j \psi''_k \psi_j \psi_k^* \psi'_j \psi_i^* \\ &= \sum_{i,j,m,k,l} \psi'_m (\psi''_i \psi_j \psi_i^* \psi_j^*) \psi_m^* (\psi'_k \psi_l \psi_k^* \psi_l^*) \\ &= \text{id}_{\rho'} \otimes \tilde{c}(\rho, \rho'') \circ \tilde{c}(\rho, \rho') \otimes \text{id}_{\rho''}, \end{aligned}$$

which is one of the braid relations. One easily sees that $\tilde{c}(\rho, \rho') \tilde{c}(\rho', \rho) = \mathbf{1}$, thus $\tilde{c}(\cdot, \cdot)$ is a symmetry for the tensor category Δ_f . If ρ and ρ' are irreducible and localized spacelike to each other, the normal commutation relations of the corresponding fields imply that $\tilde{c}(\rho, \rho) = \pm 1$, where the minus sign occurs iff ρ and ρ'' are fermionic. Now, for irreducible ρ, ρ' define $c(\rho, \rho') = \pm c(\rho, \rho')$, where we take the minus sign iff ρ and ρ' are fermionic, and extend c to reducible objects by naturality. Then $c(\rho, \rho') = 1$ whenever ρ, ρ' are localized spacelike. Now it follows from the uniqueness result Proposition 200 that $E(\varepsilon(\rho, \rho')) = c(\rho, \rho')$. Thus $E_{\mathfrak{F}}$ is a symmetric tensor functor in the sense that it maps the symmetry ε of Δ_f to the symmetry c of the category \mathcal{H} of Hilbert spaces. Equivalently, E is a symmetric tensor functor into the category of super Hilbert spaces equipped with the symmetry \tilde{c} . ■

Thus every complete normal field net \mathfrak{F} gives rise to a strict symmetric $*$ -preserving fiber functor $E_{\mathfrak{F}}$. Denoting by $\mathfrak{F}_{E_{\mathfrak{F}}}^R$ the Roberts field net associated to the latter, our aim is to construct an isomorphism $\mathfrak{F} \cong \mathfrak{F}_{E_{\mathfrak{F}}}^R$.

THEOREM 281. *Let \mathfrak{F} be a complete normal field net for \mathfrak{A} and $E_{\mathfrak{F}} : \Delta_f \rightarrow \mathcal{SH}$ the fiber functor from Proposition 280. Then there is a unitary equivalence $\mathfrak{F}_{E_{\mathfrak{F}}} \rightarrow \mathfrak{F}$ of field nets.*

Proof. By Proposition 280, there is a symmetric $*$ -preserving fiber functor $E_{\mathfrak{F}} : \Delta_f \rightarrow \mathcal{SH}$. By the concrete Tannaka theorem (Theorem 377 of the appendix), the compact group $G_{E_{\mathfrak{F}}}$ of unitary monoidal natural transformations of E is unitarily represented on the spaces $E_{\mathfrak{F}}(\rho)$. On the other hand, the compact group G coming with our field net \mathfrak{F} also acts on these spaces, providing a homomorphism $G \rightarrow G_{E_{\mathfrak{F}}}$. This homomorphism is injective since G is concretely given as a group of unitaries on the Hilbert space H where \mathfrak{F} lives. It is also surjective, since otherwise $\pi|\mathfrak{A}$ would contain representations that are not in Δ_f , contradicting the assumption that \mathfrak{F} is a complete field net. Thus the given group G can be identified with the one reconstructed from the fiber functor $E_{\mathfrak{F}}$. For every $\sigma \in \hat{F}$ we define a projection \mathcal{E}_σ on \mathfrak{F} as in Definition 274. We denote by \mathfrak{F}_0 the algebraic direct sum $\bigoplus_{\sigma \in \hat{G}} \mathcal{E}_\sigma(\mathfrak{F})$, which is the same as $\{F \in \mathfrak{F} \mid \mathcal{E}_\sigma(F) = 0 \text{ for almost all } \sigma \in \hat{G}\}$.

We now define a map $\gamma : \mathfrak{F}_{E_{\mathfrak{F}},0}^R \rightarrow \mathfrak{F}$ by $\gamma : (A, \rho, \psi) \mapsto \pi_0(A)\psi$. At first sight, this formula looks strange, but it makes perfect sense since $\psi \in E_{\mathfrak{F}}(\rho)$,

where $E_{\mathfrak{F}}(\rho)$ by definition is a subspace of \mathfrak{F} . As usual, γ is well defined since, for $T \in \text{Hom}(\rho, \rho')$,

$$\gamma((AT, \rho, \psi)) = \pi_0(AT)\psi = \pi_0(A)E_{\mathfrak{F}}(T)\psi = \gamma((A, \rho', E_{\mathfrak{F}}(T)\psi)).$$

Furthermore,

$$\begin{aligned} \gamma((A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2)) &= \gamma((A_1\rho_1(A_2), \rho_1\rho_2, \psi_1 \otimes \psi_2)) = \pi_0(A_1\rho_1(A_2))\psi_1\psi_2 \\ &= \pi_0(A_1)\psi_1\pi_0(A_2)\psi_2 = \gamma((A_1, \rho_1, \psi_1))\gamma((A_2, \rho_2, \psi_2)), \end{aligned}$$

where we have used $\psi_1 \in E_{\mathfrak{F}}(\rho_1) = \{F \in \mathfrak{F} \mid F\pi_0(A) = \pi_0(\rho(A))F\}$. Thus γ is an algebra homomorphism. This, together with $(A, \rho, \psi) = (A, \iota, 1)(\mathbf{1}, \rho, \psi)$ implies that γ is a $*$ -homomorphism provided $\gamma(F^*) = \gamma(F)^*$ for $F = (\mathbf{1}, \rho, \psi)$. Now, using the $*$ -operation on \mathfrak{F}^R defined in Proposition 263, we have

$$\gamma((\mathbf{1}, \rho, \psi)^*) = \gamma((R^*, \bar{\rho}, (\mathcal{J}E(\bar{R}^*))\psi)) = \pi_0(R^*)(\mathcal{J}E(\bar{R}^*))\psi.$$

On the other hand, $\gamma((\mathbf{1}, \rho, \psi))^* = \psi^*$, thus γ is a $*$ -homomorphism provided $\psi^* = R^*(\mathcal{J}E(\bar{R}^*))\psi$ holds for all $\psi \in E(\rho)$.

Now, for any $\bar{\psi} \in E(\bar{\rho})$, we have $R^*\bar{\psi}\rho(A) = R^*\bar{\rho}\rho(A)\bar{\psi} = AR^*\bar{\psi}$, thus $(R^*\bar{\psi})^* \in E(\rho)$. Applying this to $\bar{\psi} = \mathcal{J}E(\bar{R}^*)\psi \in E(\bar{\rho})$, we see that $\psi^* = R^*(\mathcal{J}E(\bar{R}^*))\psi$ holds iff $\psi^*\psi' = R^*(\mathcal{J}E(\bar{R}^*))\psi\psi'$ for all $\psi' \in E(\rho)$.

By Proposition 344 of the Appendix, $(E(\rho), E(R), E(\bar{R}))$ is a conjugate of $E(\rho)$ in the category of Hilbert spaces. (Or super Hilbert spaces. This doesn't matter since we don't use the symmetry.) Thus there are bases $\{e_i\}, \{f_i\}$ of $E(\rho)$ and $E(\bar{\rho})$, respectively, with dual bases $\{\hat{e}_i\}, \{\hat{f}_i\}$ in $\widehat{E(\rho)}, \widehat{E(\bar{\rho})}$ such that

$$E(R) = \sum_i f_i \otimes e_i, \quad E(\bar{R}) = \sum_i e_i \otimes f_i, \quad E(R)^* = \sum_i \hat{f}_i \otimes \hat{e}_i, \quad E(\bar{R})^* = \sum_i \hat{e}_i \otimes \hat{f}_i.$$

Thus, for $\psi \in E(\rho), \bar{\psi} \in E(\bar{\rho})$, we have

$$\langle \mathcal{J}E(\bar{R}^*)\psi, \bar{\psi} \rangle = \left(\sum_i \hat{e}_i \otimes \hat{f}_i \right) (\psi \otimes \bar{\psi}) = \sum_i \hat{e}_i(\psi)\hat{f}_i(\bar{\psi})$$

and therefore $\mathcal{J}E(\bar{R}^*)\psi = \sum_i \widehat{\hat{e}_i(\psi)} f_i$. Thus

$$E(R)^*((\mathcal{J}E(\bar{R}^*)\psi) \otimes \psi') = \left(\sum_i \hat{f}_i \otimes \hat{e}_i \right) \left(\sum_j \widehat{\hat{e}_j(\psi)} f_j \otimes \psi' \right) = \sum_i \widehat{\hat{e}_i(\psi)} \hat{e}_i(\psi') = \langle \psi, \psi' \rangle.$$

Now, in \mathfrak{F} , the left hand side equals $R^*(\mathcal{J}E(\bar{R}^*)\psi)\psi'$ and the right hand side equals $\psi^*\psi'$, proving the desired identity $\psi^* = R^*(\mathcal{J}E(\bar{R}^*))\psi$.

Now, for $(A, \rho, \psi) \in \mathcal{F}_{E_{\mathfrak{F}}, 0}^R$ is is clear that $\gamma((A, \rho, \psi))$ is contained in a finite dimensional G -stable subspace of \mathfrak{F} and thus in \mathfrak{F}_0 . Every $F \in \mathfrak{F}_0$ is a sum of finitely many terms of the form $\mathcal{E}_\sigma(F)$ with $\sigma \in \widehat{G}$. Picking an irreducible

subspace H_σ of isometries in \mathfrak{F}_0 transforming according to the class σ , there is an endomorphism $\rho \in \Delta_f$ induced by the subspace H_σ . Since every $F \in \mathcal{E}_\sigma(\mathfrak{F})$ is a linear combination $\sum_i A_i \psi_i$ with $A_i \in \mathfrak{A}, \psi_i \in H_\sigma$, we have $F = \gamma(\sum_i (A, \rho, \psi_i))$, proving $\gamma(\mathfrak{F}_{E_{\mathfrak{F}},0}^R) = \mathfrak{F}_0$.

Let $(A, \rho, \psi) \in \mathfrak{F}_0^R$. By construction of \mathfrak{F}_0 , we have a finite sum representation $(A, \rho, \psi) = \sum_i (A_i, \rho_i, \psi_i)$, where the ρ_i are irreducible and mutually non-isomorphic. Now $\gamma((A, \rho, \psi)) = \sum_i A_i \psi_i$, where the spaces $E(\rho_i) \subset \mathfrak{F}$ transform under mutually inequivalent irreducible representations of G . Thus $\gamma((A, \rho, \psi)) = 0$ iff $A_i \psi_i = 0$ for all i . Isometries transforming according to a representation in the class σ . Since by harmonic analysis, every $F \in \mathfrak{F}$ has a unique representation of the form $F = \sum_\sigma A_{\sigma,i} \psi_i^\sigma$, this implies that for each i we have $(A_i, \rho_i, \psi_i) = 0$. Thus γ is injective.

We have thus proven that $\gamma : \mathfrak{F}_{E_{\mathfrak{F}},0}^R \rightarrow \mathfrak{F}_0$ is an isomorphism. Since the vacuum state $\omega_0^{\mathfrak{F}} = (\Omega_0, \cdot \Omega)$ of \mathfrak{F} is by assumption gauge invariant, the states $\omega_0^{\mathfrak{F}} \circ \gamma$ and $\omega_0^{\mathfrak{A}} \circ m$ on $\mathfrak{F}_{E_{\mathfrak{F}},0}^R$ coincide, implying that the completed nets are unitarily equivalent in their GNS representations. ■

COROLLARY 282. *Every complete normal field net \mathfrak{F} is unitarily equivalent to a Roberts field net \mathfrak{F}_E^R , where it doesn't matter which fiber functor E we use.*

NOTE 283. As promised, we return to the issue of strictness of the functor $F : \mathcal{C} \rightarrow \mathcal{SH}_f$ that was assumed in the construction of the field net, but not proven in the appendix. In the latter, we constructed a non-strict fiber functor, i.e. a functor $E : \mathcal{C} \rightarrow \mathcal{SH}_f$ together with natural isomorphisms $d_{\rho,\rho'}^E : E(\rho) \otimes E(\rho') \rightarrow E(\rho \otimes \rho')$ and $e^E : \mathbf{1}_{\mathcal{SH}} \rightarrow E(\iota_\Delta)$ satisfying Eqns. (56), (57). The construction of the (algebraic) field algebra \mathfrak{F}_0 in Subsection 10.2 can easily be generalized to this situation: The product of fields is defined by

$$(A_1, \rho_1, \psi_1)(A_2, \rho_2, \psi_2) := \left(A_1 \rho_1(A_2), \rho_1 \otimes \rho_2, d_{\rho_1, \rho_2}^E(\psi_1 \otimes \psi_2) \right)$$

and the unit is $(\mathbf{1}, \iota, e^E \mathbf{1}_{\mathcal{C}})$. Now associativity and the unit property are obvious consequences of Eqns. (56), (57). The rest of the constructions and proofs goes through as before, just carrying the unitaries d^E, e^E along. An interesting consequence of this and of Proposition 280 is that we can prove the existence of a strict fiber functor $E' : \Delta_f \rightarrow \mathcal{SH}'$, where \mathcal{SH}'_f is a strictification of the category of finite dimensional super Hilbert spaces. This is consistent with strictification results in category theory. (Strictification of tensor categories is nicely treated in [Kassel, 1995, Chap. XI], but for strictification of tensor functors the best reference remains [Joyal and Street, 1993a, Sect. 1].)

10.6 Further relations between \mathfrak{A} and \mathfrak{F} , and a Galois interpretation

In Section 9 we have discussed at length the structure of the superselection sectors of a net \mathfrak{A} of observables in relation to the harmonic analysis of the action of a (global) gauge group on a field net \mathfrak{F} . Note that we did not claim that all

DHR representations of the fixed point net $\mathfrak{A} = \mathfrak{F}^G$ are connected to the vacuum representation by the fields in \mathfrak{F} . In order to see that this is in general false, consider a theory \mathfrak{A} with non-trivial DHR-category and take $\mathfrak{F} := \mathfrak{A}$ as ‘field net’, acted upon by the trivial group $G = \{e\}$. Obviously, all DHR representations of \mathfrak{A} are not created by the action of \mathfrak{F} on H_0 . In the special case where \mathfrak{F} is Bosonic and itself satisfies all the requirements on an observable net, it may have non-trivial DHR sectors. Restricting a DHR representation π of \mathfrak{F} with $d(\pi) < \infty$ to \mathfrak{A} , one obtains a DHR representation of \mathfrak{A} of the same dimension, which therefore decomposes into a finite direct sum of irreducibles. If π is irreducible and inequivalent to the vacuum representation π_0 of \mathfrak{F} , then all the irreducible representations of \mathfrak{A} obtained in this way are disjoint from those contained in $\pi_0|_{\mathfrak{A}}$. We refrain from a further analysis of this issue. We do, however, wish to point out that one can specify conditions on a net \mathfrak{F} implying that all DHR representations of \mathfrak{A} are contained in $\pi_0|_{\mathfrak{A}}$. This involves the net-cohomology or local 1-cohomology developed by J.E. Roberts and reviewed, e.g., in [Roberts, 1990, §3.4]. We refrain from even attempting to give precise statements and only say the following: If \mathfrak{F} has ‘quasi-trivial 1-cohomology’ and is acted upon by a compact group G of global gauge symmetries, then the equivalent (by Proposition 206) categories $DHR_f(\mathfrak{A}) \simeq \Delta_f(\mathfrak{A})$ are equivalent, as symmetric tensor categories to $\text{Rep}_f G$. In [Buchholz *et al.*, 1992] it is shown, e.g., that the theory of a free massive scalar field has quasi-trivial 1-cohomology. Thus, if one takes \mathfrak{F} to be the direct product of N copies of such fields (of the same mass) then $SO(N)$ acts on \mathfrak{F} . Therefore, $\Delta_f(\mathfrak{A}) \simeq \text{Rep}_f G$ whenever $G \subset SO(N)$ is a closed subgroup and $\mathfrak{A} = \mathfrak{F}^G$. In [Doplicher and Piacitelli, 2002] this observation is combined with a limit construction to prove that every (second countable) compact group arises as a DHR gauge group. In a similar fashion, one shows that if \mathfrak{F} is the theory of a massive Fermion with its canonical $\mathbb{Z}/2$ -symmetry, then $\Delta_f(\mathfrak{F}^{\mathbb{Z}/2}) \simeq \text{Rep}_f \mathbb{Z}/2$.

There are results in the opposite direction, i.e. from the superselection structure of \mathfrak{A} to that of \mathfrak{F} . By [Doplicher and Roberts, 1990, Theorem 3.6], which we have not covered entirely in Section 9, the field net reconstructed in [Doplicher and Roberts, 1990] and in Section 10 above satisfies ‘twisted Haag duality’. In particular, if \mathfrak{A} has no Fermionic representations then \mathfrak{F} satisfies Haag duality. In this case, one can study the categories $DHR(\mathfrak{F})$ or $\Delta(\mathfrak{F})$. In [Conti *et al.*, 2001], the following has been proven:

THEOREM 284. *Let \mathfrak{A} be a net of observables such that there are at most countably many DHR representations of finite dimension, all of which are Bosonic. Then the complete field net \mathfrak{F} has no non-trivial DHR representations of finite dimension.*

Rather than trying to comment on the many other known results related to those treated in the preceding sections, we close this section by commenting on a very satisfactory *mathematical* interpretation of DHR/DR theory. We are referring to the evident analogy between this theory and the Galois theory of algebraic field extensions. (It should be clear that in the latter context, by ‘field’

we mean the algebraic structure of which $\mathbb{Q}, \mathbb{R}, \mathbb{C}$ are examples, not the theory of classical or quantum fields.) A field \mathbb{F} in the latter sense is called algebraically closed if every polynomial $P(x)$ with coefficients in \mathbb{F} has a zero in \mathbb{F} . (P then is a product of linear factors $x - a$.) Every field \mathbb{F} is a subfield of an essentially unique algebraically closed field $\overline{\mathbb{F}}$ that is an algebraic extension of \mathbb{F} . The latter means that $\overline{\mathbb{F}}$ is obtained by adjoining, typically transfinitely, solutions of polynomial equations to \mathbb{F} . The group $G_{\mathbb{F}} = \text{Aut}_{\mathbb{F}}(\overline{\mathbb{F}})$ is compact, and one has a bijective correspondence between intermediate fields $\mathbb{F}' \subset \overline{\mathbb{F}}, \mathbb{F}' \supset \mathbb{F}$ and closed subgroups $H \subset G$. (The correspondence is given by $H \mapsto \overline{\mathbb{F}}^H, \mathbb{F}' \mapsto \text{Aut}_{\mathbb{F}'}(\overline{\mathbb{F}})$.) A similar Galois correspondence holds in AQFT, cf. e.g. [Conti *et al.*, 2001; Carpi and Conti, 2001]. In view of Theorem 284, the construction of the complete DR field net is entirely analogous to that of the algebraic closure and can be considered as the passage to a simpler or better behaved theory. Conversely, just as taking the fixed field of an algebraically closed field \mathbb{F} under the action of a closed subgroup $G \subset \text{Aut}_{\mathbb{F}}$ will result in an algebraically non-closed field \mathbb{F}^G , taking the G -fixed subnet of a net \mathfrak{F} with trivial category $\Delta_f(\mathfrak{F})$ (more precisely, quasi-trivial 1-cohomology) will result in a net with non-trivial category $\Delta_f(\mathfrak{A})$. Thus the ‘complication’ manifested by a non-trivial DHR-category $\Delta_f(\mathfrak{A})$ indicates that the theory \mathfrak{A} ‘really’ is just a subtheory of a simpler one.

Physically, however, it is not at all clear whether the ‘observable’ net $\mathfrak{A} = \mathfrak{F}^G$ with its non-trivial representation category $\Delta_f(\mathfrak{A})$ or the ‘field net’ \mathfrak{F} with trivial $\Delta_f(\mathfrak{F})$ but non-trivial global symmetry group G is more fundamental – at least when \mathfrak{F} is Bosonic. In [Haag, 1996] it is argued that the ‘right’ description of the physical situation should be in terms of a net without any global symmetries. (On the other hand, in [Haag, 1996, Section III.4.2] one finds a tentative postulate on a ‘good’ net \mathfrak{A} of observables that implies triviality of $\Delta_f(\mathfrak{A})$. As the above discussion shows, it will be very hard to find a theory that has both a trivial DHR category Δ_f and trivial global symmetry group! The theory of a single free massive Bose field is one of the rare examples.) Whether or not one subscribes to these views, from a mathematical point of view, both nets \mathfrak{A} and \mathfrak{F} contain the same amount of information. This equivalence is in fact a useful tool, since it permits to view many problems from different angles. For example, while a spin statistics theorem can be proven in a ‘field’ framework, its physical interpretation may be clearer in the ‘observable’ setting.

10.7 Spontaneous symmetry breaking

So far, our entire analysis has presupposed the axiom of Haag duality for the theory \mathfrak{A} . Haag duality played an important rôle in our analysis of the category $\Delta(\mathfrak{A})$, but is needed also to establish the equivalence between the $\Delta(\mathfrak{A})$ and the representations satisfying the a priori physically motivated DHR criterion (Definition 204). Thus, while it seems that the study of DHR representations is physically motivated also for non-Haag dual nets, our mathematical analysis soon gets stuck. We will therefore briefly comment on an approach to resolve this issue, which turns

out to have a profound physical interpretation.

DEFINITION 285. Let $O \mapsto \mathfrak{R}(O)$ be a net of von Neumann algebras on a vacuum Hilbert space H_0 . The *dual net* \mathfrak{R}^d of \mathfrak{R} is the assignment $O \mapsto \mathfrak{R}(O)'$.

If we have $O_1 \subset O_2$ then $O_2' \subset O_1'$, thus $\mathfrak{R}(O_2') \subset \mathfrak{R}(O_1')$, and therefore $\mathfrak{R}^d(O_1) \subset \mathfrak{R}^d(O_2)$. Thus the dual net really satisfies isotony. Microcausality of \mathfrak{R} is equivalent to $\mathfrak{R}(O) \subset \mathfrak{R}(O')' = \mathfrak{R}^d(O)$, or briefly $\mathfrak{R} \subset \mathfrak{R}^d$, and Haag duality of \mathfrak{R} is equivalent to $\mathfrak{R} = \mathfrak{R}^d$. If $\mathfrak{A}_1 \subset \mathfrak{A}_2$ (in the sense of an inclusion for every O) then $\mathfrak{A}_2^d \subset \mathfrak{A}_1^d$, thus $\mathfrak{A} \subset \mathfrak{A}^{dd}$, and a standard argument shows $\mathfrak{A}^d = \mathfrak{A}^{ddd}$. Note, however, that microcausality of \mathfrak{R} does not imply microcausality of \mathfrak{R}^d ! This motivates the following

DEFINITION 286. A net $O \mapsto \mathfrak{R}(O) \subset B(H_0)$ satisfies *essential duality* if both \mathfrak{R} and the dual net \mathfrak{R}^d (both indexed by double cones) satisfy microcausality.

LEMMA 287. *If \mathfrak{R} satisfies essential duality then $\mathfrak{R}^d = \mathfrak{R}^{dd}$, i.e. \mathfrak{R}^d satisfies Haag duality.*

NOTE 288. Essential duality can be difficult to verify. Fortunately, essential duality follows from *wedge duality*, to wit $\mathfrak{R}(W')' = \mathfrak{R}(W)$ for all wedge regions (the regions obtained from the standard wedge $W_0 = \{x \in \mathbb{R}^{1+s} \mid x_0 \geq |x_1|\}$ by Poincaré transformations). Besides being much easier to verify than essential duality, wedge duality is a very fundamental property that might well be required of any ‘reasonable’ QFT.

Assuming that \mathfrak{R} satisfies essential duality, \mathfrak{R}^d satisfies Haag duality and the D(H)R analysis applies to it. Thus we obtain a symmetric tensor *-category with conjugates $\Delta_f^d(\mathfrak{R}) := \Delta_f(\mathfrak{R}^d) \simeq DHR_f(\mathfrak{R}^d)$, and we can construct the complete DR field net \mathfrak{F} associated with $(\mathfrak{R}^d, \Delta_f(\mathfrak{R}^d))$. One thus has an inclusion $\mathfrak{R} \subset \mathfrak{R}^d \subset \mathfrak{F}$ of nets. The DR gauge group acts on \mathfrak{F} and we have $\mathfrak{F}^G = \mathfrak{R}^d$ and also $G = \text{Aut}_{\mathfrak{R}^d}(\mathfrak{F})$. Since the group G is implemented by unitaries that leave the vacuum vector fixed, G consists of ‘unbroken symmetries’. One can now define a larger group

$$\widehat{G} = \text{Aut}_{\mathfrak{R}}(\mathfrak{F})$$

and topologize it suitably. Now $G \subset \widehat{G}$ consists precisely of the elements of \widehat{G} that are unitarily implemented. The point is that the net \mathfrak{R} acts irreducibly on H_0 , thus a unitary whose adjoint action leaves all algebras $\mathfrak{R}(O)$ pointwise fixed must be a multiple of the identity also on \mathfrak{R}^d .

Concerning the categories associated with \mathfrak{R} , little can be said about the category $\Delta(\mathfrak{R})$, but Roberts proved the existence of an extension functor $K : DHR(\mathfrak{R}) \rightarrow DHR(\mathfrak{R}^d)$ such that $K(\pi)|_{\mathfrak{R}} = \pi$ for every $\pi \in DHR(\mathfrak{R})$, cf. [Roberts, 1990, §3.4]. (Again, a crucial rôle is played by the theory of local 1-cohomology. Furthermore, this result breaks down in less than three spacetime dimensions due to the phenomenon of solitonic representations.) This functor actually is an equivalence, thus spontaneous symmetry breakdown doesn’t manifest itself in the superselection structure.

For the detailed analysis we refer to [Roberts, 1974; Roberts, 1990] and the remarkable paper [Buchholz *et al.*, 1992], in which the Goldstone phenomenon is analyzed in the context of algebraic QFT.

Notes: DHR superselection theory originates in a four-paper series: [Doplicher *et al.*, 1969a] starts with a field algebra and gauge group and then derives properties of the superselection sectors. [Doplicher *et al.*, 1969b] reconstructs the field algebra and gauge group from the category of representations in the special case where the objects are all one dimensional (i.e. the equivalence classes of objects of Δ_f form an abelian group with the monoidal product and conjugates as inverses). [Doplicher *et al.*, 1971] defines the symmetry $\varepsilon_{\rho_1, \rho_2}$, and uses it to give the statistical classification of objects of Δ .

For surveys of DHR theory in general, see [Roberts, 1970; Roberts, 1990; Roberts, 2004; Fredenhagen, 1992; Fredenhagen, 1994], [Araki, 1999, Ch. 6], and [Haag, 1996, Ch. IV.2].

The full proof of the DR reconstruction theorem is distributed over [Doplicher *et al.*, 1974; Doplicher and Roberts, 1972; Doplicher and Roberts, 1989] and [Doplicher and Roberts, 1990]. The alternative approach to the reconstruction theorem that we use in this paper is based on [Roberts, ND] and [Deligne, 1990], incorporating simplifications of the latter due to [Bichon, 1998] and ourselves.

For informal expositions of the DR reconstruction theorem, see [Doplicher and Roberts, 1987; Doplicher, 1991; Doplicher, 1992; Doplicher, 1993; Doplicher, 1995]. For an interesting description of the goal of reconstructing fields and a gauge group, written before a solution was obtained, see [Roberts, 1975].

11 FOUNDATIONAL IMPLICATIONS OF THE RECONSTRUCTION THEOREM

We now return to the foundational questions (Section 7) that motivated our investigation. We also point out a few other cases where discussions in the philosophical literature (e.g. about permutation symmetry and the identity of particles) might benefit from the study of superselection theory.

11.1 Algebraic Imperialism and Hilbert Space Conservatism

DHR superselection theory sheds light on some questions that philosophers have asked about the role of inequivalent representations of the algebra of observables. But it will not answer all of our questions. We first bracket off those questions for which DHR theory provides little help, and then we proceed to develop a case for the relevance of DHR theory to foundational questions.

The DHR analysis requires that we fix a vacuum state ω_0 , and hence a base representation (\mathcal{H}_0, π_0) . Inequivalent DHR representations do not correspond to different vacua; rather, they correspond to different local excitations of one and the same vacuum state. So, DHR theory effectively ignores the question of how to choose a vacuum representation. (But note that the power of the DHR analysis

strongly suggests — against the Algebraic Imperialist — that representations *are* essential components of the physical content of the theory.)

Second, in some of the most familiar cases — e.g., the free Boson field — the DHR category is trivial. That is, $DHR(\mathfrak{A}) = \{\pi_0\}$, and so $\mathfrak{F} = \mathfrak{A}$. In this case, the vacuum representation is the *only* DHR representation (relative to itself). Thus, in such cases, the elaborate apparatus of DHR superselection theory seems to provide little insight into the physical importance of inequivalent representations. (However, if we are able to find a physical reason for choosing a preferred vacuum representation, then the DHR analysis suggests that no other representations are relevant for explaining the phenomena.)

Finally, even in cases where $DHR(\mathfrak{A})$ is nontrivial, the field algebra itself has inequivalent representations. (After all, it's just another large C^* -algebra.) And one might worry that the same Conservative versus Imperialist debate will arise with respect to the field algebra.

But DHR theory has something to say about inequivalent DHR representations, and about representations of the field algebra. First, the field algebra \mathfrak{F} is constructed concretely as operators on a Hilbert space \mathcal{H} ; i.e. \mathfrak{F} comes with a preferred representation. (Recall that the preferred representation of \mathfrak{F} is on a Hilbert space \mathcal{H} that decomposes into a direct sum of superselection sectors for \mathfrak{A} .) Of course, we could consider other representations of \mathfrak{F} . But in other representations of \mathfrak{F} , we no longer have the intuitive interpretation of elements of \mathfrak{F} as intertwiners between the DHR sectors of \mathfrak{A} . If the physically meaningful quantities have to have some connection to observable phenomena, then the interpretation of elements of \mathfrak{F} in terms of \mathfrak{A} might be thought to be a necessary condition for interpretation; and so the given representation of \mathfrak{F} might be preferred for these reasons.

So, DHR theory suggests that the issue of inequivalent representations does not come up for the field algebra. Regarding the issue of inequivalent representation of the observable algebra, we can divide the problem into several questions:

1. Is there a physically privileged vacuum state/representation? What features pick it out?
2. Are all physical representations in the DHR category of some vacuum state? (We are ignoring for the time being theories with long range forces (see [Buchholz and Fredenhagen, 1982]). In more general settings, we expect the *form* of the question to remain the same: do we require physical states to be reachable from a fixed vacuum state by the action of an appropriate set of fields?)
3. If the answer to the second question is No, then how should we compare representations that are not connected to a vacuum representation by fields to representations that are?

Let's suppose that the first question has been answered in the affirmative, and that the vacuum representation (\mathcal{H}_0, π_0) is fixed. Suppose also that $DHR(\mathfrak{A})$ is

nontrivial. Then how should we think about the inequivalent representations in $DHR(\mathfrak{A})$? A naive transcription of Hilbert Space Conservatism to the current context would tell us the following: the representations in $DHR(\mathfrak{A})$ are analogous to competing theories: one is correct, and the others are surplus structure. The naive transcription of Algebraic Imperialism to the current context would say: the representations in $DHR(\mathfrak{A})$ are surplus structure; the physical content of the theory is in \mathfrak{A} , the abstract algebra of observables.

Both Conservatism and Imperialism are based on an oversimplified view of the formalism: it is supposed that the elements of reality correspond to operators in the abstract algebra or in some Hilbert space representation thereof, and that the possible states are either all states on the abstract algebra or some particular folium of states. But the fundamental insight of DHR theory is that the set of representations itself has structure, and it is *this* structure that explains phenomena. So, a more adequate position would take all the representations more seriously. Hence, we propose that according to *Representation Realism*, the content of the theory is given by: (i) the net $O \mapsto \mathfrak{A}(O)$, (ii) the dynamics on the quasilocal algebra (i.e. the representation of the translation group in $\text{Aut}\mathfrak{A}$), and (iii) the symmetric tensor $*$ -category $DHR(\mathfrak{A})$ of DHR representations.

Recall that the Conservative claims to have the advantage in terms of explanatory power: more structure (provided by choosing a representation) provides more elements of reality, and so more satisfying explanations. But DHR superselection theory shows that this claimed advantage is misleading: the focus on one representation ignores the most interesting structure, namely the relations between representations. Indeed, if we committed ourselves to one representation, and ignored others, we would have no field operators, no gauge group, no definition of Bose and Fermi fields, no definition of antiparticles, etc..

And yet there is a strong *prima facie* objection to Representation Realism: since the Hamiltonian is always an observable, no possible dynamical evolution can take us from a state in one representation to a state in an inequivalent representation. So, inequivalent representations are dynamically isolated from each other, and abstract relations between them cannot explain the features of states in the representation that best describes our universe.

The fact that the Hamiltonian is an observable — hence cannot map states from one sector to states in another — raises a further problem for our interpretation of field operators. Recall that we speak of “creating” states in a sector \mathcal{H}_ρ by acting on the vacuum with elements from the field algebra. That is, we can choose $F \in H_\rho \subseteq \mathfrak{F}$ such that $F\Omega \in \mathcal{H}_\rho$, where $(\mathcal{H}_\rho, \pi_\rho)$ is disjoint from the vacuum representation (\mathcal{H}_0, π_0) . The talk of “creation” here suggests that we are talking about some sort of dynamical process. On the one hand, $F \in \mathfrak{F}$ can be chosen unitary, so structurally the map $\Omega \mapsto F\Omega$ looks like dynamics. But since the Hamiltonian is an observable, the transition $\Omega \mapsto F\Omega$ is not dynamically allowable. So, in what sense are states in \mathcal{H}_ρ accessible from the vacuum? Is the key insight behind superselection rules that there are two notions of dynamic accessibility? If so, then how are we to understand the differences between these two notions?

11.2 Explanatory relations between representations

If we consider a C^* -algebra \mathfrak{A} with no further structure, then the mathematically definable relations between representations (and hence, between states in the folia of these representations) are exhausted by the following table:

π_1 and π_2 are equivalent
π_1 and π_2 are quasiequivalent
π_1 and π_2 are disjoint
π_1 and π_2 are weakly equivalent

Table 1. Relations Between Representations of \mathfrak{A}

Outside of the fourth relation (which makes special reference to the topology of the state space), these relations are precisely those that can be defined in an arbitrary $*$ -category \mathcal{C} with subobjects. Two objects X, Y in \mathcal{C} are equivalent if there is a unitary $u \in \text{Hom}(X, Y)$; are quasiequivalent if there is an isometry $v \in \text{Hom}(X, Y)$; and are disjoint just in case they are not quasiequivalent.

Consider now the normal state space K of a C^* -algebra \mathfrak{A} . The GNS theorem provides a map $\omega \mapsto (\mathcal{H}_\omega, \pi_\omega)$ from K into the representation category of \mathfrak{A} . We then use this map to induce relations corresponding to those in Table 1 on K : we talk about equivalent, quasiequivalent, and disjoint states. Furthermore, the individual folia (sets of states whose GNS representations are quasiequivalent) have a rich geometrical structure which corresponds exactly to the normal state space of $\mathfrak{B}(\mathcal{H})$ for some Hilbert space \mathcal{H} . Thus, within a folium we have a notion of “transition probability” between pure states (corresponding to rays in \mathcal{H}), and a three place relation “ ω is a superposition of ω_1 and ω_2 .” However, if two states lie in disjoint folia, then these relations trivialize. The transition probability between disjoint states is zero, and no state can be a superposition of states from a different folia. It seems that the only physically suggestive thing we can say about states from different folia is that they are “orthogonal.”

It is precisely the preceding considerations that have lead philosophers to worry about inequivalent representations. The worry is based on the fact that disjoint representations seem to be *competitors*, and nothing more. In order to alleviate worries about the competition between representations, some philosophers [Clifton and Halvorson, 2001a; Halvorson, 2004] go so far as to claim that these representations are “complementary” descriptions of the phenomena (in the sense of Bohr). The word “complementarity” is of course meant to suggest that the representations are not *merely* competitors, and the choice of one does not need to be seen as completely ruling out the relevance of another.

We wish to replace suggestive — and possibly misleading — terminology with some real facts about the relationships between inequivalent representations. To illustrate what we mean by this, consider the case of *group* representations: let $\text{Rep}_f G$ be the category of unitary representations of a compact group G . $\text{Rep}_f G$ is

not only a $*$ -category, but it has a monoidal product and conjugates. That is, for objects X, Y in $\text{Rep}_f G$, there is a product object $X \otimes Y$, and a conjugate object \overline{X} . For our purposes, this is the crucial difference between group representations and the representations of an arbitrary C^* -algebra \mathfrak{A} . For an arbitrary C^* -algebra \mathfrak{A} , there is no product of representations, or conjugate of a representation.

In the case of compact group representations, typically $X \in \text{Rep}_f G$ will be disjoint from both $X \otimes Y$ and \overline{X} . But in this case, we are not tempted to see X is *merely* as a competitor of $X \otimes Y$, or of \overline{X} ; there are some interesting relations between these representations. Roughly speaking, information about X gives us information about $X \otimes Y$ and \overline{X} . Thus, although these representations are technically “disjoint,” they are not completely unrelated to each other.⁵

One of the main accomplishments of the DHR analysis and DR reconstruction theorem is to show that the category Δ_f of physical representations is a tensor $*$ -category with conjugates; indeed the Embedding Theorem (see the Appendix) shows that Δ_f is equivalent to the category $\text{Rep}_f G$ for some compact group G . The obvious question then is whether these additional relations on the category of representations can help us get past the idea that disjoint representations are merely competing descriptions.

An analogy to states might be helpful. Consider a pair H_1, H_2 of Hilbert spaces, and let $\psi_i \in H_i$ be unit vectors. Now consider the following two “descriptions of reality”:

1. The state is ψ_1 .
2. The state is $\psi_1 \otimes \psi_2$.

What do we say here: are these competing descriptions? In one sense, (1) and (2) are competitors, because they cannot both be fully adequate descriptions of reality at the same time. However, (1) and (2) are not competitors in the same sense that, say, two orthogonal vectors in a single Hilbert space are competitors. The two state descriptions are not merely competitors, because there is an interesting sense in which ψ_1 is a “part” of $\psi_1 \otimes \psi_2$. Indeed, information about ψ_1 (e.g. expectation values it assigns to observables) does in fact give us information about $\psi_1 \otimes \psi_2$ because of the canonical mappings between H_1 and $H_1 \otimes H_2$.

Now let π_1, π_2 be objects in the DHR category Δ_f , and suppose (as will often be the case) that the representations π_1 and $\pi_1 \otimes \pi_2$ are disjoint. Are these competing descriptions? Again, π_1 and $\pi_1 \otimes \pi_2$ are competitors in the sense that if the state of an object (or of the universe?) is in $\mathcal{H}_{(\pi_1 \otimes \pi_2)}$ then it is not in \mathcal{H}_{π_1} . Nonetheless, π_1 and $\pi_1 \otimes \pi_2$ are not merely competitors, because in one sense π_1 is “part” of $\pi_1 \otimes \pi_2$.

⁵But note also: Philosophers of physics have so far not worried about inequivalent *group* representations as competing descriptions of reality. And for good reason, because group elements are not observables, and groups do not have states. Another insight of DHR theory is to show that physicist’s intuitions about group representations are not totally baseless, because in fact the interesting (DHR) representations of the observable algebra correspond to representations of a compact group.

But two words of caution should be issued here. First, we must be cautious with the use of the “part” metaphor. For example, Δ_f can have a nontrivial representation π such that $\pi \otimes \pi$ is equivalent to the vacuum representation. Then it is not so clear that we should say that “ π is part of $\pi \otimes \pi$.” Second, there is one significant disanalogy between the case of states ψ_1 and $\psi_1 \otimes \psi_2$ and the case of representations π_1 and $\pi_1 \otimes \pi_2$: the two representations are GNS representations of states on a *single* C^* -algebra \mathfrak{A} . Hence we can directly compare the expectation values that these states assign to observables in \mathfrak{A} , and they will disagree significantly (indeed, for any $\varepsilon > 0$ there is an observable $A \in \mathfrak{A}$ such that $\|A\| \leq 1$ and $\|\omega_1(A) - \omega_2(A)\| > 2 - \varepsilon$). Thus, there is a clear, empirically verifiable sense in which states in π_1 are competitors with states in $\pi_1 \otimes \pi_2$.

Finally, there is an interesting physical relation between a DHR representation π and its conjugate $\bar{\pi}$, even though π and $\bar{\pi}$ are typically disjoint. In short, $\bar{\pi}$ is like an inverse of π : if π is irreducible, then $\bar{\pi}$ is the unique irreducible representation such that $\pi \otimes \bar{\pi}$ contains a copy of the vacuum representation. In fact, when $\pi = \pi_0 \circ \rho$ where ρ is a dimension 1 element of Δ_f , $d(\rho) = 1$, then this is the exact relation: ρ is an automorphism and $\bar{\rho} = \rho^{-1}$. In terms of field operators, if F creates the charge ξ , then \bar{F} annihilates the charge ξ . Furthermore, when π admits a particle interpretation, then the states in the folium of $\bar{\pi}$ are the antiparticle states of the states in the folium of π [Doplicher *et al.*, 1969b].

11.3 Fields as theoretical entities, as surplus structure

From the standpoint of superselection theory, there is a sharp distinction between observable and unobservable fields, namely, a field operator is an observable iff it is invariant under all gauge transformations. To what extent does this distinction between fields and observables match up with the philosopher of science’s distinction between theoretical and observational components of a theory? Even if the two notions are not exactly the same, the connection is suggestive. In particular, it seems interesting to ask about the extent to which the field plus gauge part of QFT is fixed by the observable algebra.

First, the notion of equivalent systems of field operators seems a fairly close analogue of the philosopher’s notion of “theoretical equivalence.”

DEFINITION 289. Let $\mathfrak{F}_1 = (\mathfrak{F}_1, \mathcal{H}_1, G_1)$ and $\mathfrak{F}_2 = (\mathfrak{F}_2, \mathcal{H}_2, G_2)$ be local field systems with gauge symmetry for (\mathfrak{A}, ω) . (See Defn. 247 on p. 808.) Then \mathfrak{F}_1 and \mathfrak{F}_2 are *theoretically equivalent* just in case they are unitarily equivalent as systems of local field operators. (See Defn. 10 on p. 816.)

REMARK 290. (i) This definition is not fully adequate, because it does not make reference to dynamics. For example, this definition entails that the free Bose field nets for different positive masses are theoretically equivalent. For a fully adequate definition, it would probably be necessary to require that the unitary mapping $W : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ also intertwines the dynamical groups on the two Hilbert spaces. (ii) If \mathfrak{F}_1 and \mathfrak{F}_2 are theoretically equivalent, then they are equivalent in *all* physically relevant senses (modulo dynamics): they have the same type of commutation

relations (either both have normal or abnormal commutation relations), they have isomorphic gauge groups, etc..

The working analogy also suggests that we define “observational equivalence” between two theories in terms of some equivalence between their nets of observable algebras. There are a myriad number of ways we could explicate the notion of observational equivalence in this setting; philosophers have their work cut out for them. The following two definitions give extremely weak notions of observational equivalence that do not take into account a representation of the algebra of observables.

DEFINITION 291. Let \mathfrak{F}_1 and \mathfrak{F}_2 be two local field systems with gauge symmetry, and let \mathfrak{A}_1 and \mathfrak{A}_2 be the fixed point algebras; i.e.,

$$\mathfrak{A}_i = \{A \in \mathfrak{F}_i : \alpha_g(A) = A, \text{ for all } g \in G_i\}.$$

Then we say that \mathfrak{F}_1 and \mathfrak{F}_2 are *weakly observationally equivalent* just in case there is a $*$ -isomorphism α from the algebra \mathfrak{A}_1 onto the algebra \mathfrak{A}_2 .

DEFINITION 292. Let \mathfrak{F}_1 and \mathfrak{F}_2 be two local field systems with gauge symmetry, and let \mathfrak{A}_1 and \mathfrak{A}_2 be their fixed point nets; i.e. for each double cone O ,

$$\mathfrak{A}_i(O) = \{A \in \mathfrak{F}_i(O) : \alpha_g(A) = A, \text{ for all } g \in G_i\}.$$

Then we say that \mathfrak{F}_1 and \mathfrak{F}_2 are *observationally equivalent* just in case there is a net isomorphism $\alpha : \mathfrak{A}_1 \rightarrow \mathfrak{A}_2$ (see Defn. 118 on p. 767).

REMARK 293. The first definition is weaker because it does not require that the net structure be preserved by the $*$ -isomorphism α .

Again, the definitions omit reference to dynamics, which would be an important component of a fully adequate treatment of observational equivalence. Nonetheless, even with these definitions, we can make some sense of remarks about underdetermination of fields by observables, or about the physical equivalence of different field theories.

1. (Construction of observationally equivalent theories) The DR reconstruction theorem provides a general, nontrivial recipe for constructing non-equivalent theories that are observationally equivalent: If (\mathfrak{A}, ω_0) has nontrivial DHR superselection sectors, then it can be embedded into two nonequivalent field algebras \mathfrak{F}_1 and \mathfrak{F}_2 . Indeed, \mathfrak{A} is always a field algebra over itself (but incomplete), but the field algebra \mathfrak{F} from the DR reconstruction theorem is complete.
2. (Elimination of parafields) It has long been thought that parafields are a theoretical artifact. More precisely, it has been claimed that every parafield theory is “physically equivalent” to a theory with normal commutation relations (see [Araki, 1961]). The DR reconstruction theorem partially validates this claim by showing that every parafield theory is *observationally equivalent* to a theory with normal commutation relations. Indeed, suppose that

\mathfrak{F}_1 is a parafield theory. Then we can extract the observable algebra \mathfrak{A} contained in \mathfrak{F}_1 , and apply the DR reconstruction theorem to construct a field algebra \mathfrak{F}_2 with normal commutation relations. Since \mathfrak{F}_1 and \mathfrak{F}_2 have the same net of local observable algebras, they are observationally equivalent.

3. Some have claimed that the relation between quantum fields (the field algebra \mathfrak{F}) and observables (the observable algebra \mathfrak{A}) is analogous to relation between coordinates and a manifold in differential geometry. However, the DR reconstruction theorem shows that (subject to normal commutation relations), there is a *unique* field net \mathfrak{F} and gauge group G compatible with the observable algebra (\mathfrak{A}, ω_0) . Thus, there is a strong disanalogy between the two cases, since there seems to be no sense in which one coordinate system of a manifold is a better representation of reality than another coordinate system.

Finally, we are in a position to get a clear picture of the interrelations between questions about inequivalent representations and questions about gauge invariance.

According to a common interpretive stance, if two states of a physical system can be identified under a gauge symmetry, then those two states are different descriptions of the same underlying reality. So, for the purposes of counting states, we should look to the quotient of the state space under gauge orbits. Similarly, it is thought that the “real” quantities of a theory are the gauge invariant quantities (see [Earman, 2004]).

In the setting of DHR superselection theory, the algebra of observables \mathfrak{A} is precisely the gauge invariant part of the field algebra \mathfrak{F} , that is,

$$\mathfrak{A} = \{A \in \mathfrak{F} : \alpha_g(A) = A, \text{ for all } g \in G\},$$

where G is the gauge group. This of course means that for any observable $A \in \mathfrak{A}$, there is no difference between a state ψ and the gauge transformed state $U(g)\psi$. (Of course, if ψ is a state vector in the vacuum representation, then $U(g)\psi = \psi$, since the representation of the gauge group there is trivial.) So, if the common interpretive stance is correct, the physical content of the theory is in the observable algebra \mathfrak{A} ; the fields are “descriptive fluff.”

So suppose that we ignore the field algebra \mathfrak{F} , and just look to the observable algebra \mathfrak{A} to provide the physical content of the theory. But what should we say about the representations of \mathfrak{A} ? Are they just descriptive fluff? If not, then is there one correct representation of \mathfrak{A} , or do we somehow need inequivalent representations in order to account for all of the physical facts?

The DR reconstruction theorem shows that the preceding two sets of questions — regarding the status of gauge variant quantities on the one hand, and representations on the other hand — are tightly intertwined. The full structure of the theory, field algebra \mathfrak{F} and gauge group G , is uniquely recoverable (modulo completeness, and normal commutation relations) from the structure of the category $DHR(\mathfrak{A})$ of representations. The ontological significance of the gauge variant fields

is closely aligned with the ontological significance of inequivalent representations. (We will revisit this question in the next section when we discuss permutation symmetry.)

Of course, there is a crucial disanalogy between the global gauge symmetries in DHR superselection theory, and the local gauge symmetries of electromagnetism or general relativity. But it is not clear that this disanalogy makes the DR reconstruction theorem any less interesting for understanding the relation between gauge symmetry and superselection rules.

11.4 *Statistics, permutation symmetry, and identical particles*

Philosophers have taken an active interest in the differences between the Maxwell-Boltzmann statistics of classical physics, and the Bose-Fermi statistics of quantum physics. Indeed, it has been provocatively claimed that Bose-Fermi statistics is explained by permutation invariance — i.e. there are no physical differences between states with permuted particle labels — and that this entails that quantum particles are not “individuals” in the same sense that classical particles are. (See [French, 2000; French and Rickles, 2003] for an account of the argument.)

But such discussions can be hampered by an overly simplistic formalism. In particular, states are usually identified with unit vectors (or at best with rays) in a single Hilbert space, and no account is given of the status of non-permutation invariant operators. It might be helpful then to look at the issue from the perspective of a more flexible formalism that allows a distinction between fields and observables, and in which states can be represented by vectors in inequivalent representations of the algebra of observables.

There is another reason why the issue of permutation invariance should be revisited within the context of QFT. Some of the literature suggests that the metaphysical problem about the individuality of particles is *aufgehoben* in the transition to QFT, because: (i) QFT is about fields and not particles, and (ii) the Fock space formalism of QFT already identifies permuted states, and so rules out a notion of particles as individuals. We have already noted that it is not so clear how to make sense of the idea that QFT is about fields as opposed to particles. Furthermore, the DR reconstruction theorem shows precisely how to make sense of non-permutation invariant quantities and states in a theory that is manifestly permutation invariant.

It is not surprising that DHR theory is relevant for the issue of permutation invariance and statistics: one of the original goals of DHR was to clarify the role of statistics in QFT. Riding the crest of a wave of mathematical success, Roberts made the following bold claim about the DHR analysis of statistics:

One of the insights provided by the study of superselection sectors concerns the origin of what is termed the ‘statistics’ of a particle. . . . Now just as a theory should determine its particle states so should it determine the statistics of these particles. Ordinary quantum mechanics ignores this challenge saying in effect that the statistics of particles is

one of the parameters determining the theory, the one telling you what symmetry the n -particle states have. QFT does a little better: it says that the statistics of the particles in the theory is determined by the commutation relations of fields at spacelike separations. ... In adopting the philosophy that the local observables determine the theory, we are forced to meet this challenge in full. [Roberts, 1976, p. 203]

In the remainder of the paper from which the quote is taken, Roberts shows how Bose-Fermi particle statistics emerges naturally from the DHR analysis of physical representations of the algebra of observables.

Roberts' claim is of crucial relevance to the philosophical debate about statistics and identical particles. The philosophers have asked, "what explains Bose-Fermi statistics?" Roberts' answer is that the explanation comes from the structure of the category of representations of the observable algebra.

Let us recall then how the Bose-Fermi distinction is supposed to emerge from the DHR analysis. In Section 8.3, an object ρ of the category Δ is shown to have an intrinsic dimension $d(\rho)$. The dimension is finite iff ρ has a conjugate; in this case we define a unitary operator $\Theta_\rho \in \text{End}(\rho)$ called the *twist* of ρ . If ρ is irreducible, then $\Theta_\rho = \omega_\rho \text{id}_\rho$ where $\omega_\rho = \pm 1$. We then stipulate that a "Bosonic" object is one with $\omega_\rho = 1$ and a "Fermionic" object is one with $\omega_\rho = -1$.

Of course, ρ is not the sort of thing that the philosophers have been calling Bosonic or Fermionic — it is not a wavefunction. To connect the two pieces of formalism, recall that an object of Δ_f (endomorphisms of the algebra of observables) corresponds to a representation $\pi_0 \circ \rho$ of the algebra of observables. So, we call the representation $\pi_0 \circ \rho$ Bosonic when ρ is Bosonic, and Fermionic when ρ is Fermionic. Finally, we then call a state ("wavefunction") Bosonic if it is in the folium of a Bosonic representation, and Fermionic if it is in the folium of a Fermionic representation. The claims to be assessed are: (i) does this stipulative definition adequately reproduce the distinction between Bosonic and Fermionic wavefunctions made in elementary nonrelativistic QM; and if so, (ii) what does this tell us about permutation invariance?

The Bose-Fermi distinction in nonrelativistic QM

In nonrelativistic QM, the state space of n identical particles is the tensor product $H \otimes \cdots \otimes H$ of n copies of a Hilbert space H . The Hilbert space $H \otimes \cdots \otimes H$ is spanned by *product states*, i.e. states of the form

$$\psi_1 \otimes \cdots \otimes \psi_n,$$

with $\psi_1, \dots, \psi_n \in H$.

DEFINITION 294. We define the natural action of the permutation group S_n on $H \otimes \cdots \otimes H$ as follows. Let $\{\psi_1, \dots, \psi_m\}$ be an orthonormal basis for H , and define for each permutation σ ,

$$U(\sigma)(\psi_{i_1} \otimes \cdots \otimes \psi_{i_m}) = \psi_{\sigma(i_1)} \otimes \cdots \otimes \psi_{\sigma(i_m)},$$

and extend $U(\sigma)$ by linearity to all of H .

If $\dim H > 1$, then the representation U of S_n is reducible. It contains copies of the two one-dimensional representations of S_n , namely the identity representation $S_n \rightarrow \mathbf{1}$, and the alternating representation. The subspace of vectors of $H \otimes \cdots \otimes H$ transforming according to the identity representation is called the *symmetric* subspace, and the subspace of vectors transforming according to the alternating representation is called the *antisymmetric* subspace. Vectors in the symmetric subspace are called Bosonic, and vectors in the antisymmetric subspace are called Fermionic. These traditional definitions have served as the basis of discussions of permutation invariance in the philosophical literature.

In QM, states of n particles that differ only by permuting labels — for example, $\psi_1 \otimes \psi_2$ versus $\psi_2 \otimes \psi_1$ — should not be counted separately. For the purposes of statistical weighting, these two symbols represent one state. This has been stated as the principle of Permutation Invariance.

Permutation Invariance (PI): Let \mathfrak{A} be the observables for the n particle system. Then for each state ψ , and for each permutation $\sigma \in S_n$, we have

$$\langle U(\sigma)\psi, AU(\sigma)\psi \rangle = \langle \psi, A\psi \rangle.$$

Permutation Invariance is sometimes also called the *Indistinguishability Postulate*: two states that differ by a permutation are indistinguishable (i.e. there is no measurement that can distinguish between the two). It has been claimed that PI entails that state has Bose or Fermi statistics, no states with “parastatistics” are allowed.

Dichotomy: For each state vector ψ and permutation σ , we have $U(\sigma)\psi = \pm\psi$.

(See [van Fraassen, 1991, pp. 389ff] for an account of attempts to prove Dichotomy from PI. See [Butterfield, 1993] for further details.) In other words, the states that are not in the symmetric or antisymmetric subspaces are surplus structure.

This leaves us with a number of puzzles. First, what do we say about the vectors in $H \otimes \cdots \otimes H$ that are not in the symmetric or antisymmetric subspaces? Are they surplus structure? Are they possibilities that are contingently not realized in nature? More generally, not all vectors in $H \otimes \cdots \otimes H$ are of a definite symmetry type; and even among those that are of a definite symmetry type, not all are totally symmetric or totally antisymmetric. For any irreducible representation ξ of S_n we say that a wavefunction ψ in $H \otimes \cdots \otimes H$ is of symmetry type ξ just in case ψ is contained in the subspace corresponding to the representation ξ . Then $H \otimes \cdots \otimes H$ is the direct sum of subspaces of definite symmetry type vectors. But now the principle of plenitude suggests that there should be particles of every symmetry type. Why do we not see them?

An intrinsic characterization of symmetric and antisymmetric subspaces

We began with the full n -particle Hilbert space $H \otimes \cdots \otimes H$, and then we reduced to the symmetric and antisymmetric subspaces. We were then left wondering what to do with the remaining elements of $H \otimes \cdots \otimes H$.

The intrinsic description of the symmetric and antisymmetric subspaces is that they are representations of the group S_n . (In fact, they are quasiequivalent to the one dimensional irreducible representations of S_n .) So we can also work backwards. That is, if we are given a representation (H, π) of S_n , we can ask after the intrinsic characterization of this representation. Recall that the irreducible representations of S_n are in one-to-one correspondence with Young tableaux with n boxes (see [Simon, 1996]). There is a natural grouping of representations of S_n into para-Bose and para-Fermi categories: we specify the representation (H, π) by a pair of numbers $(d(\pi), \omega_\pi)$, with $d(\pi) \in \{1, 2, \dots, n\}$ and $\omega_\pi = \pm 1$.

1. For $(d, +1)$, all Young tableaux whose columns have length less than or equal d . (In this case, we say that π has *Para-Bose statistics* of order d .)
2. For $(d, -1)$, all Young tableaux whose rows have length less than or equal d . (In this case, we say that π has *para-Fermi statistics* of order d .)

Clearly the division of representations into para-Bose and para-Fermi is mutually exclusive, but not exhaustive. (e.g., there are representations of S_n that contain copies of both the **1** representation and the alternating representation.)

Now suppose that we are in the following position (described vividly by Roberts in the opening quote): we are given a pure state ω of the algebra of observables \mathfrak{A} and we asked whether its “intrinsic” statistics of its states is Bosonic or Fermionic. What can we do? First we construct the GNS representation (\mathcal{H}, π) induced by ω . At least this makes things more concrete. But the Hilbert space \mathcal{H} is not itself a tensor product, and so there is no natural representation of S_n on \mathcal{H} . Nor would it help to construct tensor products of elements of \mathcal{H} in the normal way, because for $\psi \otimes \cdots \otimes \psi$ is trivially Bosonic. So, the obvious approach does not seem to tell us anything about the *intrinsic* symmetry type of elements of \mathcal{H} .

The key insight again comes from the DHR analysis: the representation (\mathcal{H}, π) is naturally isomorphic to an object ρ of a symmetric tensor $*$ -category, namely the category Δ_f of localized transportable endomorphisms. Since Δ_f has products we can construct $\rho \otimes \rho$, and the symmetry $\varepsilon_{\rho, \rho}$ gives us notion of permuting $\rho \otimes \rho$. [Recall that $\varepsilon_{\rho, \rho} \in \text{Hom}(\rho \otimes \rho)$.] As we will see in the following section, this gives us a natural representation u of S_n in $\text{End}(\rho \otimes \rho)$. Furthermore, the pair $(d(\rho), \omega_\rho)$, where $d(\rho)$ is the dimension of ρ , and ω_ρ is the statistical phase of ρ , coincide with the classification of u as a para-Bose or para-Fermi representation of S_n . We will also see that this natural representation u of S_n corresponds to a permutation of wavefunctions in the “larger” Hilbert space of the field algebra \mathfrak{F} .

Representations of S_n in a symmetric tensor $$ -category*

Unitary representations of the permutation group S_n arise naturally in tensor $*$ -categories with a unitary symmetry. Let $(\mathcal{C}, \otimes, \mathbf{1})$ be a tensor $*$ -category with unitary symmetry $c_{X,Y}$. Fix an object $X \in \text{Obj}(\mathcal{C})$, and define a map $u : S_2 \rightarrow \text{End}(X \otimes X)$ by setting

$$u((1)) = \text{id}_{X \otimes X}, \quad u((1, 2)) = c_{X,X}.$$

Since $(c_{X,X})^2 = \text{id}_{X \otimes X}$, u is a unitary representation of S_2 in $\text{End}(X \otimes X)$. This construction can be iterated: define a map $u : S_n \rightarrow \text{End}(X \otimes \cdots \otimes X)$ by setting

$$u((i, i + 1)) = \text{id}_X \otimes \cdots \otimes c_{X,X} \otimes \cdots \otimes \text{id}_X.$$

It is easy to verify that u extends uniquely to a unitary representation of S_n in $\text{End}(X \otimes \cdots \otimes X)$.

FACT 295. Let \mathcal{C} be a tensor $*$ -category with unitary symmetry and conjugates. Then for each irreducible object $X \in \text{Obj}(\mathcal{C})$ the induced unitary representation u of S_n in $\text{End}(X \otimes \cdots \otimes X)$ is para-Bose of order $d(X)$ if $\omega_X = +1$, and is para-Fermi of order $d(X)$ if $\omega_X = -1$. Furthermore, the statistical phase ω_X is the trace of $u((1, 2)) = c_{X,X}$. (See Appendix B for more details.)

The physical interpretation becomes more clear in the presence of field operators. (Of course, the point of the Reconstruction Theorem is to show that such field operators are always available.) Let $(\mathcal{H}, \mathfrak{F}, (G, k))$ be a field system with gauge symmetry for the observable algebra \mathfrak{A} and vacuum state ω . Let O_1, \dots, O_n be mutually spacelike separated regions. Let ρ be an irreducible object in Δ_f . Then using the transportability of ρ we can choose $F_i \in \mathfrak{F}(O_i)$ such that $F_i \Omega$ is in the sector $\mathcal{H}_{\hat{\rho}}$. (Recall that sectors are labeled by unitary equivalent classes $\hat{\rho}$ of objects in Δ_f .) In other words, F_i creates the charge $\hat{\rho}$ in the region O_i . Let σ be a permutation of $\{1, \dots, n\}$ and consider the following two state vectors in \mathcal{H} :

$$(52) \quad \psi_1 \times \psi_2 \times \cdots \times \psi_n \equiv F_1 F_2 \cdots F_n \Omega,$$

$$(53) \quad \psi_{\sigma(1)} \times \psi_{\sigma(2)} \times \cdots \times \psi_{\sigma(n)} \equiv F_{\sigma(1)} F_{\sigma(2)} \cdots F_{\sigma(n)} \Omega.$$

These two vectors are typically distinct. In fact, if the field net has normal commutation relations then we can compute the difference between the two vectors. Supposing that σ only permutes two numbers, the two vectors will be the same if ρ is Bosonic, and the two vectors will differ by a sign if ρ is Fermionic. However, the two vectors always induce the same state on the algebra of observables $\pi(\mathfrak{A})$. Indeed, if $\rho_i \in \Delta_f(O_i)$ are the corresponding morphisms, then the states induced by the two vectors, respectively, are

$$(54) \quad \omega \circ (\rho_1 \otimes \rho_2 \otimes \cdots \otimes \rho_n) = \omega \circ (\rho_1 \rho_2 \cdots \rho_n),$$

$$(55) \quad \omega \circ (\rho_{\sigma(1)} \otimes \rho_{\sigma(2)} \otimes \cdots \otimes \rho_{\sigma(n)}) = \omega \circ (\rho_{\sigma(1)} \rho_{\sigma(2)} \cdots \rho_{\sigma(n)}).$$

Since endomorphisms that are localized in spacelike separated regions commute, these two states are equal. Thus, permutation invariance holds for the observables, but not for the fields.

The interpretive upshot of the the DHR treatment of statistics is as follows: permutation invariance is a gauge symmetry in the sense that it leaves all observables' values invariant, but changes the values assigned to field operators. Are two states related by a permutation the same or different? Of course, the answer to the mathematical question is clear: the states of the observable algebra are the same, the states of the field algebra are different. So, whether or not we take permutations to correspond to “real” changes depends on what we think about the status of the field operators. So the issue of permutation invariance is just a special version of the issue of gauge invariance, and accordingly is tightly connected to the question of the status of inequivalent representations.

Parastatistics and nonabelian gauge groups

The abstract Tannaka Theorem (Appendix B) shows that each symmetric tensor $*$ -category (STC $*$) \mathcal{C} is equivalent to the representation category $\text{Rep}_f(G, k)$ of a compact supergroup (G, k) . Applied to our current topic, the theorem shows that the category Δ_f of localized transportable morphisms is equivalent to the representation category of the gauge group. Furthermore, Section B.9 shows that each object X of an STC $*$ gives rise naturally to a unitary representation of the symmetric group S_n in $\text{End}(X \otimes \cdots \otimes X)$, and this representation corresponds to the intrinsic statistical characterization of X . Now, we know that the categorical dimension of a representation (H, π) of (G, k) corresponds to the dimension of the underlying Hilbert space H . Hence:

LEMMA 296. *The category $\text{Rep}_f(G, k)$ has irreducible objects of dimension greater than 1 iff G is nonabelian.*

Sketch of proof. The set of irreducible representations of G separates the elements of G . Hence for $gh \neq hg$, there is an irreducible representation (H, π) such that $\pi(g)\pi(h) \neq \pi(h)\pi(g)$. Therefore $\dim H \geq 2$. ■

It immediately follows from Fact 295, in conjunction with the fact that the embedding functor preserves dimension, that:

PROPOSITION 297. *There is an irreducible object X of $\mathcal{C} \simeq \text{Rep}_f(G, k)$ with parastatistics iff the corresponding group G is nonabelian.*

Applied to our current case, this means that there are representations and states with parastatistics iff the gauge group G is nonabelian.⁶ But we have good reasons to think that the case of nonabelian gauge groups is physically relevant. So, the DHR approach ignores worries about the supposed nonexistence of paraparticle states, and undermines claims that there is a proof of Dichotomy.

⁶But there is an ambiguity in “parastatistics.” We mean para-Bose or para-Fermi statistics, not mixtures of Bose and Fermi statistics.

Braid group statistics

Recall from Section 8.1 that when spacetime is dimension 2, then $\varepsilon_{\rho_1, \rho_2}$ is not necessarily a symmetry on Δ_f , but only a braiding. In this case, objects in Δ_f are not classified according to representations of the symmetric group S_n ; rather, objects in Δ_f are classified in terms of representations of the braid group B_n . In physical terms, states might not be permutation invariant, but satisfy the more general braid group statistics.

DEFINITION 298. The *braid group* B_n on n strands is the group generated by the set $\{\sigma_1, \dots, \sigma_{n-1}\}$ satisfying the equations

- (1) $\sigma_i \sigma_j = \sigma_j \sigma_i, \quad |i - j| \geq 2,$
- (2) $\sigma_{i+1} \sigma_i \sigma_{i+1} = \sigma_i \sigma_{i+1} \sigma_i.$

The braid group on n strands can be given the following heuristic, geometric description: a braid consists of two sets $\{a_1, \dots, a_n\}$ and $\{b_1, \dots, b_n\}$ of base points, and n strands of yarn, where each yarn strand has one end attached to a_i , and the other end attached to b_j , and each base point touches only one strand of yarn. We identify two configurations if one can be transformed into the other without disconnecting the strands from the base points. In this picture, the identity braid has a_i connected to b_i , and no twists. The generating element σ_i can be thought of as the simple braid where: a_i is connected to b_{i+1} , a_{i+1} is connected to b_i , and these two strands are twisted only once. (Otherwise, the σ_i braid looks just like the identity braid.) Under this interpretation, the product gh of two braids is the braid formed by attaching the ending points of g to the beginning points of h .

PROPOSITION 299. For each $n \in \mathbb{N}$, the mapping

$$\varepsilon_\rho^{(n)}(\sigma_i) = \rho^{i-1}(\varepsilon_\rho) = I_\rho \times \dots \times I_\rho \times \varepsilon_\rho \times I_\rho \times \dots \times I_\rho,$$

defines a unitary representation of the braid group B_n in $\text{End}(\rho \otimes \dots \otimes \rho)$. For each $i, j \in \mathbb{N}$ with $i \leq j$ we have

$$\varepsilon_\rho^{(i)}(g) = \varepsilon_\rho^{(j)}(\varphi_{ij}(g)), \quad \forall g \in S_i.$$

(A proof of this Proposition can be found in [Rehren *et al.*, 1990]. Alternatively, it is obvious given the considerations in Appendix B.) In other words, the product object $\rho \otimes \dots \otimes \rho$ carries a unitary representation of the braid group, which is induced by the unitary operators of the form:

$$I_\rho \times \dots \times \varepsilon_{\rho, \rho} \times \dots \times I_\rho.$$

This represents an elementary permutation of the i -th and $(i + 1)$ -st copy of ρ .

There is a natural homomorphism of the braid group onto the symmetric group. This is obvious when we recall that the definition of the symmetric group is exactly

the same as the definition of the braid group with the additional condition that each generator is its own inverse. Hence, van Dyck's theorem [Hungerford, 1980, p. 78] entails that the obvious map f on generators extends uniquely to a group homomorphism $\bar{f} : B_n \rightarrow S_n$. So, each representation π of S_n yields a representation $\pi \circ \bar{f}$ of B_n . In slogan form: a system that obeys permutation statistics also obeys braid statistics.

Recall now the worrisome argument for the existence of paraparticles: There should be particles corresponding to all irreducible representations of S_n . For $n \geq 3$, there are non Bose or Fermi representations of S_n , so there should be paraparticles.

Now we can see that either something is wrong with this argument, or the problem is much more severe than we thought. Since any system that has S_n as a symmetry group also has B_n as a symmetry group, the argument commits us to predicting the existence of particles corresponding to all irreducible representations of B_n . But B_n has infinitely many irreducible representations. (Indeed, its representations have so far resisted classification.) Furthermore, we could now repeat the argument for *any* group K that can be mapped homomorphically onto B_n , and there is an infinite chain of larger groups with this property. Therefore, the principle of plenitude applied to group representations predicts more particles than we could ever possibly describe.

Notes: For discussions of statistics of DHR representations, see [Roberts, 1976; Doplicher, 1975].

ACKNOWLEDGEMENTS

HH wishes to thank: Michael Müger for teaching him about the Doplicher–Roberts Theorem; the editors for their helpful feedback and patience; and David Baker, Tracy Lupher, and David Malament for corrections. MM gratefully acknowledges financial support by NWO via the Pioneer grant no.616.062.384 of N. P. Landsman and thanks Julien Bichon for a critical reading of the appendix and useful comments.

BIBLIOGRAPHY

- [Araki, 1961] Huzihiro Araki. On the connection of spin and commutation relations between different fields. *Journal of Mathematical Physics*, 2:267–270, 1961.
- [Araki, 1963] Huzihiro Araki. A lattice of von Neumann algebras associated with the quantum theory of a free Bose field. *Journal of Mathematical Physics*, 4, 1963.
- [Araki, 1999] Huzihiro Araki. *Mathematical theory of quantum fields*. Oxford University Press, New York, 1999.
- [Arntzenius, 2003] Frank Arntzenius. Is quantum mechanics pointless? *Philosophy of Science*, 70(5):1447–1457, 2003.
- [Bacciagaluppi, 1994] G. Bacciagaluppi. Separation theorems and Bell inequalities in algebraic quantum mechanics. In *Symposium on the Foundations of Modern Physics 1993: Quantum*

- Measurement, Irreversibility and the Physics of Information*, pages 29–37. World Scientific, Singapore, 1994.
- [Baez *et al.*, 1992] John Baez, Irving Segal, and Zheng-Fang Zhou. *Introduction to algebraic and constructive quantum field theory*. Princeton University Press, Princeton, NJ, 1992.
- [Baez, 1987] John Baez. Bell's inequality for C^* -algebras. *Letters in Mathematical Physics*, 13(2):135–136, 1987.
- [Baumgärtel and Wollenberg, 1992] Hellmut Baumgärtel and Manfred Wollenberg. *Causal nets of operator algebras*. Akademie-Verlag, Berlin, 1992.
- [Bell, 1987] J. S. Bell. Beables for quantum field theory. In *Quantum implications*, pages 227–234. Routledge & Kegan Paul, London, 1987.
- [Beltrametti and Cassinelli, 1981] Enrico G. Beltrametti and Gianni Cassinelli. *The logic of quantum mechanics*. Addison-Wesley Publishing Co., Reading, Mass., 1981.
- [Bichon, 1998] Julien Bichon. Trivialisations dans les catégories tannakiennes. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 39(4):243–270, 1998.
- [Bichon, 1999] Julien Bichon. Galois extension for a compact quantum group. math.QA/9902031, 1999.
- [Bichon, ND] Julien Bichon. Trivialisations dans les catégories tannakiennes. Preprint version of [Bichon, 1998], ND.
- [Bohm and Hiley, 1995] D. Bohm and B. J. Hiley. *The undivided universe*. Routledge, London, 1995.
- [Bohr and Rosenfeld, 1950] Niels Bohr and Leon Rosenfeld. Field and charge measurements in quantum electrodynamics. *Physical Review*, 78:794–798, 1950.
- [Borchers, 1967] H.-J. Borchers. A remark on a theorem of B. Misra. *Communications in Mathematical Physics*, 4:315–323, 1967.
- [Borchers, 2000] H. J. Borchers. On revolutionizing quantum field theory with Tomita's modular theory. *Journal of Mathematical Physics*, 41:3604–3673, 2000.
- [Bostelmann, 2000] Henning Bostelmann. *Lokale Algebren und Operatorprodukte am Punkt*. PhD thesis, Universität Göttingen, 2000.
- [Bostelmann, 2004] Henning Bostelmann. Phase space properties and the short distance structure in quantum field theory. math-ph/0409070, 2004.
- [Bratteli and Robinson, 1997] Ola Bratteli and Derek W. Robinson. *Operator algebras and quantum statistical mechanics*, volume 2. Springer, NY, 1997.
- [Buchholz and Fredenhagen, 1982] Detlev Buchholz and Klaus Fredenhagen. Locality and the structure of particle states. *Communications in Mathematical Physics*, 84:1–54, 1982.
- [Buchholz and Summers, ND] Detlev Buchholz and Stephen Summers. Scattering theory in relativistic quantum field theory. In J. P. Francoise, G. Naber, and T. S. Tsun, editors, *Encyclopedia of Mathematical Physics*. Elsevier, ND.
- [Buchholz and Verch, 1995] Detlev Buchholz and Rainer Verch. Scaling algebras and renormalization group in algebraic quantum field theory. *Reviews in Mathematical Physics*, 7(8):1195–1239, 1995.
- [Buchholz and Wichmann, 1986] Detlev Buchholz and Eyvind Wichmann. Causal independence and the energy-level density of states in local quantum field theory. *Communications in mathematical physics*, 106:321–344, 1986.
- [Buchholz *et al.*, 1987] D. Buchholz, C. D'Antoni, and K. Fredenhagen. The universal structure of local algebras. *Communications in Mathematical Physics*, 111(1):123–135, 1987.
- [Buchholz *et al.*, 1991] Detlev Buchholz, Martin Poppmann, and Ulrich Stein. Dirac versus Wigner. Towards a universal particle concept in local quantum field theory. *Physics Letters. B*, 267(3):377–381, 1991.
- [Buchholz *et al.*, 1992] Detlev Buchholz, Sergio Doplicher, Roberto Longo, and John E. Roberts. A new look at Goldstone's theorem. *Reviews in Mathematical Physics*, Special Issue:49–83, 1992.
- [Buchholz, 1974] Detlev Buchholz. Product states for local algebras. *Communications in Mathematical Physics*, 36:287–304, 1974.
- [Buchholz, 1998] Detlev Buchholz. Scaling algebras in local relativistic quantum physics. In *Proceedings of the International Congress of Mathematicians, Vol. III (Berlin, 1998)*, volume Extra Vol. III, pages 109–112 (electronic), 1998.
- [Butterfield, 1993] Jeremy Butterfield. Interpretation and identity in quantum theory. *Studies in History and Philosophy of Modern Physics*, 24:443–476, 1993.

- [Carpi and Conti, 2001] Sebastiano Carpi and Roberto Conti. Classification of subsystems for local nets with trivial superselection structure. *Communications in Mathematical Physics*, 217:89–106, 2001.
- [Chaiken, 1967] Jan M. Chaiken. Finite-particle representations and states of the canonical commutation relations. *Annals of Physics*, 42:23–80, 1967.
- [Chaiken, 1968] Jan M. Chaiken. Number operators for representations of the canonical commutation relations. *Communications in Mathematical Physics*, 8:164–184, 1968.
- [Clifton and Halvorson, 2001a] Rob Clifton and Hans Halvorson. Are Rindler quanta real? Inequivalent particle concepts in quantum field theory. *British Journal for the Philosophy of Science*, 52(3):417–470, 2001.
- [Clifton and Halvorson, 2001b] Rob Clifton and Hans Halvorson. Entanglement and open systems in algebraic quantum field theory. *Studies in History and Philosophy of Modern Physics*, 32(1):1–31, 2001.
- [Clifton, 1995] Rob Clifton. Independently motivating the Kochen-Dieks modal interpretation of quantum mechanics. *British Journal for the Philosophy of Science*, 46(1):33–57, 1995.
- [Clifton, 2000] Rob Clifton. The modal interpretation of algebraic quantum field theory. *Physics Letters A*, 271(3):167–177, 2000.
- [Connes and Størmer, 1978] Alain Connes and Erling Størmer. Homogeneity of the state space of factors of type III₁. *Journal of Functional Analysis*, 28(2):187–196, 1978.
- [Connes, 1994] Alain Connes. *Noncommutative geometry*. Academic Press Inc., San Diego, CA, 1994.
- [Conti *et al.*, 2001] Roberto Conti, Sergio Doplicher, and John E. Roberts. Superselection theory for subsystems. *Communications in Mathematical Physics*, 218:263–281, 2001.
- [Cushing, 1994] James Cushing. Locality/separability: Is this necessarily a useful distinction? In D. Hull, M. Forbes, and R. Burian, editors, *Proceedings of PSA 1994*, volume I, pages 107–116. Philosophy of Science Association, East Lansing, Michigan, 1994.
- [D’Antoni, 1990] C. D’Antoni. Technical properties of the quasi-local algebra. In *The algebraic theory of superselection sectors (Palermo, 1989)*, pages 248–258. World Sci. Publishing, River Edge, NJ, 1990.
- [Deligne and Milne, 1982] P. Deligne and J. S. Milne. Tannakian categories. In *Lecture Notes in Mathematics*, volume 900, pages 101–228. Springer Verlag, 1982.
- [Deligne and Morgan, 1999] Pierre Deligne and John W. Morgan. Notes on supersymmetry (following Joseph Bernstein). In *Quantum fields and strings: a course for mathematicians, Vol. 1, 2 (Princeton, NJ, 1996/1997)*, pages 41–97. Amer. Math. Soc., Providence, RI, 1999.
- [Deligne, 1990] Pierre Deligne. Catégories tannakiennes. In *The Grothendieck Festschrift, Vol. II*, volume 87, pages 111–195. Birkhäuser Boston, Boston, MA, 1990.
- [Deligne, 2002] P. Deligne. Catégories tensorielles. *Moscow Mathematical Journal*, 2(2):227–248, 2002.
- [Dickson and Clifton, 1998] Michael Dickson and Rob Clifton. Lorentz invariance in the modal interpretation. In Dennis Dieks and Pieter Vermaas, editors, *The modal interpretation of quantum mechanics*, pages 9–47. Kluwer, 1998.
- [Dieks and Vermaas, 1998] Dennis Dieks and Pieter E. Vermaas, editors. *The modal interpretation of quantum mechanics*. Kluwer Academic Publishers, Dordrecht, 1998.
- [Dieks, 2000] Dennis Dieks. Consistent histories and relativistic invariance in the modal interpretation of quantum mechanics. *Physics Letters. A*, 265(5-6):317–325, 2000.
- [Doplicher and Longo, 1984] S. Doplicher and R. Longo. Standard and split inclusions of von Neumann algebras. *Invent. Math.*, 75(3):493–536, 1984.
- [Doplicher and Piacitelli, 2002] Sergio Doplicher and Gherardo Piacitelli. Any compact group is a gauge group. *Reviews in Mathematical Physics*, 14(7-8):873–885, 2002.
- [Doplicher and Roberts, 1972] Sergio Doplicher and John E. Roberts. Fields, statistics and non-abelian gauge groups. *Communications in Mathematical Physics*, 28:331–348, 1972.
- [Doplicher and Roberts, 1987] Sergio Doplicher and John E. Roberts. C^* -algebras and duality for compact groups: why there is a compact group of internal gauge symmetries in particle physics. In *VIIIth international congress on mathematical physics (Marseille, 1986)*, pages 489–498. World Scientific Publishing, Singapore, 1987.
- [Doplicher and Roberts, 1989] Sergio Doplicher and John E. Roberts. A new duality theory for compact groups. *Inventiones Mathematicae*, 98(1):157–218, 1989.

- [Doplicher and Roberts, 1990] Sergio Doplicher and John E. Roberts. Why there is a field algebra with a compact gauge group describing the superselection structure in particle physics. *Communications in Mathematical Physics*, 131(1):51–107, 1990.
- [Doplicher *et al.*, 1969a] Sergio Doplicher, Rudolf Haag, and John E. Roberts. Fields, observables and gauge transformations. I. *Communications in Mathematical Physics*, 13:1–23, 1969.
- [Doplicher *et al.*, 1969b] Sergio Doplicher, Rudolf Haag, and John E. Roberts. Fields, observables and gauge transformations. II. *Communications in Mathematical Physics*, 15:173–200, 1969.
- [Doplicher *et al.*, 1971] Sergio Doplicher, Rudolf Haag, and John E. Roberts. Local observables and particle statistics. I. *Communications in Mathematical Physics*, 23:199–230, 1971.
- [Doplicher *et al.*, 1974] Sergio Doplicher, Rudolf Haag, and John E. Roberts. Local observables and particle statistics. II. *Communications in Mathematical Physics*, 35:49–85, 1974.
- [Doplicher, 1975] Sergio Doplicher. The statistics of particles in local quantum theories. In *International Symposium on Mathematical Problems in Theoretical Physics (Kyoto Univ., Kyoto, 1975)*, pages 264–273. Springer, Berlin, 1975.
- [Doplicher, 1991] Sergio Doplicher. Abstract compact group duals, operator algebras and quantum field theory. In *Proceedings of the International Congress of Mathematicians, Vol. I, II (Kyoto, 1990)*, pages 1319–1333, Tokyo, 1991. Math. Soc. Japan.
- [Doplicher, 1992] Sergio Doplicher. Progress and problems in algebraic quantum field theory. In *Ideas and methods in quantum and statistical physics (Oslo, 1988)*, pages 390–404. Cambridge Univ. Press, Cambridge, 1992.
- [Doplicher, 1993] Sergio Doplicher. Operator algebras and abstract duals: progress and problems. In *Quantum and non-commutative analysis (Kyoto, 1992)*, pages 405–418. Kluwer Acad. Publ., Dordrecht, 1993.
- [Doplicher, 1995] Sergio Doplicher. Quantum field theory, categories and duality. In *Advances in dynamical systems and quantum physics (Capri, 1993)*, pages 106–116. World Sci. Publishing, River Edge, NJ, 1995.
- [Driessler, 1975] W. Driessler. Comments on lightlike translations and applications in relativistic quantum field theory. *Communications in Mathematical Physics*, 44(2):133–141, 1975.
- [Earman, 2004] John Earman. Thoroughly modern McTaggart. *Philosopher's Imprint*, 2:1–28, 2004.
- [Emch, 1972] Gérard G. Emch. *Algebraic methods in statistical mechanics and quantum field theory*. Wiley Interscience, NY, 1972.
- [Fleming, 2000] Gordon N. Fleming. Reeh-Schlieder meets Newton-Wigner. *Philosophy of Science*, 67(3, suppl.):S495–S515, 2000. PSA 1998, Part II (Kansas City, MO).
- [Florig and Summers, 1997] Martin Florig and Stephen J. Summers. On the statistical independence of algebras of observables. *Journal of Mathematical Physics*, 38(3):1318–1328, 1997.
- [Folland, 1995] Gerald B. Folland. *A course in abstract harmonic analysis*. CRC Press, Boca Raton, FL, 1995.
- [Fredenhagen and Hertel, 1981] Klaus Fredenhagen and Joachim Hertel. Local algebras of observables and pointlike localized fields. *Comm. Math. Phys.*, 80(4):555–561, 1981.
- [Fredenhagen, 1981] Klaus Fredenhagen. On the existence of anti particles. *Communications in Mathematical Physics*, 79:141–151, 1981.
- [Fredenhagen, 1992] Klaus Fredenhagen. Observables, superselection sectors and gauge groups. In *New symmetry principles in quantum field theory (Cargèse, 1991)*, volume 295, pages 177–194. Plenum, New York, 1992.
- [Fredenhagen, 1993] Klaus Fredenhagen. Global observables in local quantum physics. In *Quantum and non-commutative analysis*. Kluwer, Amsterdam, 1993.
- [Fredenhagen, 1994] Klaus Fredenhagen. Superselection sectors. <http://unith.desy.de/research/aqft/lecturenotes/>, 1994.
- [French and Rickles, 2003] Stephen French and Dean Rickles. Understanding permutation symmetry. In Katherine Brading and Elena Castellani, editors, *Symmetries in Physics*, pages 212–238. Cambridge University Press, 2003.
- [French, 2000] Steven French. Identity and individuality in quantum theory. 2000. <http://plato.stanford.edu/entries/qt-idind/>.
- [Gabriel, 1962] Pierre Gabriel. Des catégories abéliennes. *Bulletin de la Société Mathématique de France*, 90:323–448, 1962.
- [Giulini, 2003] D. Giulini. Superselection rules and symmetries. In *Decoherence and the appearance of a classical world in quantum theory*. Springer, NY, 2nd edition, 2003.

- [Guido and Longo, 1992] Daniele Guido and Roberto Longo. Relativistic invariance and charge conjugation in quantum field theory. *Communications in Mathematical Physics*, 148:521–551, 1992.
- [Haag and Swieca, 1965] R. Haag and J. A. Swieca. When does a quantum field theory describe particles? *Communications in Mathematical Physics*, 1:308–320, 1965.
- [Haag, 1996] Rudolf Haag. *Local quantum physics*. Springer-Verlag, Berlin, second edition, 1996.
- [Hái, 2002] Phùng Hồ Hái. On a theorem of Deligne on characterization of Tannakian categories. In *Arithmetic fundamental groups and noncommutative algebra (Berkeley, CA, 1999)*, volume 70 of *Proc. Sympos. Pure Math.*, pages 517–531. Amer. Math. Soc., Providence, RI, 2002.
- [Halvorson and Clifton, 2000] Hans Halvorson and Rob Clifton. Generic Bell correlation between arbitrary local algebras in quantum field theory. *Journal of Mathematical Physics*, 41(4):1711–1717, 2000.
- [Halvorson and Clifton, 2002] Hans Halvorson and Rob Clifton. No place for particles in relativistic quantum theories? *Philosophy of Science*, 69(1):1–28, 2002.
- [Halvorson, 2001] Hans Halvorson. Reeh-Schlieder defeats Newton-Wigner: on alternative localization schemes in relativistic quantum field theory. *Philosophy of Science*, 68(1):111–133, 2001.
- [Halvorson, 2004] Hans Halvorson. Complementarity of representations in quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 35(1):45–56, 2004.
- [Holland, 1995] Peter R. Holland. *The quantum theory of motion*. Cambridge University Press, Cambridge, 1995.
- [Horuzhy, 1990] S. S. Horuzhy. *Introduction to algebraic quantum field theory*. Kluwer Academic Publishers, Dordrecht, 1990.
- [Huggett, 1999] Nick Huggett. On the significance of permutation symmetry. *British Journal for the Philosophy of Science*, 50(3):325–347, 1999.
- [Huggett, 2000] Nick Huggett. Philosophical foundations of quantum field theory. *British Journal for the Philosophy of Science*, 51:617–637, 2000.
- [Hungerford, 1980] Thomas W. Hungerford. *Algebra*. Springer-Verlag, New York, 1980.
- [Joyal and Street, 1991] André Joyal and Ross Street. An introduction to Tannaka duality and quantum groups. In *Category theory (Como, 1990)*, volume 1488 of *Lecture Notes in Math.*, pages 413–492. Springer, Berlin, 1991.
- [Joyal and Street, 1993a] André Joyal and Ross Street. Braided tensor categories. *Advances in Mathematics*, 102(1):20–78, 1993.
- [Kadison and Ringrose, 1997] Richard V. Kadison and John R. Ringrose. *Fundamentals of the theory of operator algebras. Vol. I*. American Mathematical Society, Providence, RI, 1997.
- [Kassel, 1995] Christian Kassel. *Quantum groups*. Springer-Verlag, New York, 1995.
- [Kitajima, 2004] Yuichiro Kitajima. A remark on the modal interpretation of algebraic quantum field theory. *Physics Letters A*, 331(3-4):181–186, 2004.
- [Kraus, 1983] Karl Kraus. *States, effects, and operations*. Springer-Verlag, Berlin, 1983.
- [Landau, 1987] Lawrence J. Landau. On the violation of Bell’s inequality in quantum theory. *Physics Letters. A*, 120(2):54–56, 1987.
- [Landsman, 1995] N. P. Landsman. Observation and superselection in quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 26(1):45–73, 1995.
- [Longo and Roberts, 1997] Roberto Longo and John E. Roberts. A theory of dimension. *K-Theory*, 11(2):103–159, 1997.
- [Longo, 1979] Roberto Longo. Notes on algebraic invariants for noncommutative dynamical systems. *Communications in Mathematical Physics*, 69(3):195–207, 1979.
- [Longo, 1989] Roberto Longo. Index of subfactors and statistics of quantum fields, I. *Communications in Mathematical Physics*, 126:217–247, 1989.
- [Mac Lane, 1998] Saunders Mac Lane. *Categories for the working mathematician*. Springer-Verlag, New York, second edition, 1998.
- [McIntosh, 1970] Alan McIntosh. Hermitian bilinear forms which are not semibounded. *Bulletin of the American Mathematical Society*, 76:732–737, 1970.
- [Müger and Tuset, 2006] Michael Müger and Lars Tuset. Monoids, embedding functors and quantum groups. math.QA/0604065, 2006.

- [Müger *et al.*, 2004] Michael Müger, John E. Roberts, and Lars Tuset. Representations of algebraic quantum groups and reconstruction theorems for tensor categories. *Algebras and Representation Theory*, 7(5):517–573, 2004.
- [Müger, 2000] Michael Müger. Galois theory for braided tensor categories and the modular closure. *Advances in Mathematics*, 150(2):151–201, 2000.
- [Myrvold, 2002] Wayne C. Myrvold. Modal interpretations and relativity. *Foundations of Physics*, 32(11):1773–1784, 2002.
- [Napiórkowski, 1972] Kazimierz Napiórkowski. On the independence of local algebras. *Reports on Mathematical Physics*, 3(1):33–35, 1972.
- [Pedersen, 1989] Gert K. Pedersen. *Analysis now*. Springer, New York, 1989.
- [Rédei, 1998] Miklós Rédei. *Quantum logic in algebraic approach*. Kluwer Academic Publishers Group, Dordrecht, 1998.
- [Redhead, 1995a] Michael L. G. Redhead. More ado about nothing. *Foundations of Physics*, 25(1):123–137, 1995.
- [Redhead, 1995b] Michael L. G. Redhead. The vacuum in relativistic quantum field theory. In David Hull, Micky Forbes, and Richard M. Burian, editors, *PSA; proceedings of the biennial meeting of the Philosophy of Science Association*, volume 2, pages 77–87, East Lansing, MI, 1995. Philosophy of Science Association.
- [Reeh and Schlieder, 1961] H. Reeh and S. Schlieder. Bemerkungen zur Unitärsäquivalenz von Lorentzinvarianten Feldern. *Nuovo cimento (10)*, 22:1051–1068, 1961.
- [Rehberg and Wollenberg, 1986] Joachim Rehberg and Manfred Wollenberg. Quantum fields as pointlike localized objects. *Mathematische Nachrichten*, 125:259–274, 1986.
- [Rehren *et al.*, 1990] Karl-Henning Rehren, Daniel Kastler, and M. Mebkhout. Introduction to the algebraic theory of superselection sectors. In Daniel Kastler, editor, *Algebraic Theory of Superselection Sectors*, pages 113–214. World Scientific, Singapore, 1990.
- [Requardt, 1986] Manfred Requardt. Reeh-Schlieder-type density results in one- and n -body Schrödinger theory and the “unique continuation problem”. *Journal of Mathematical Physics*, 27(6):1571–1577, 1986.
- [Roberts and Roepstorff, 1968] J. E. Roberts and G. Roepstorff. Some basic concepts of algebraic quantum theory. *Communications in Mathematical Physics*, 11:321–338, 1968.
- [Roberts, 1970] John E. Roberts. The structure of sectors reached by a field algebra. In *Cargèse lectures in physics, Vol. 4*, pages 61–78. Gordon and Breach, New York, 1970.
- [Roberts, 1974] John E. Roberts. Spontaneously broken symmetries and superselection rules. 1974. unpublished preprint.
- [Roberts, 1975] John E. Roberts. Must there be a gauge group? *Trudy Mat. Inst. Steklov*, 135, 1975.
- [Roberts, 1976] John E. Roberts. Statistics and the intertwiner calculus. In *C*-algebras and their applications to statistical mechanics and quantum field theory (Proc. Internat. School of Physics “Enrico Fermi”, Course LX, Varenna, 1973)*, pages 203–225. North-Holland, Amsterdam, 1976.
- [Roberts, 1990] John E. Roberts. Lectures on algebraic quantum field theory. In Daniel Kastler, editor, *The algebraic theory of superselection sectors (Palermo, 1989)*, pages 1–112. World Scientific Publishing, River Edge, NJ, 1990.
- [Roberts, 2004] John E. Roberts. More lectures on algebraic quantum field theory. In Sergio Doplicher and Roberto Longo, editors, *Noncommutative geometry*, pages 263–342. Springer, Berlin, 2004.
- [Roberts, ND] John E. Roberts. The reconstruction problem. Unpublished Manuscript, ND.
- [Rosenberg, 2000] Alexander L. Rosenberg. The existence of fiber functors. In *The Gelfand Mathematical Seminars, 1996–1999*, Gelfand Math. Sem., pages 145–154. Birkhäuser Boston, Boston, MA, 2000.
- [Rudin, 1991] Walter Rudin. *Functional analysis*. McGraw-Hill Inc., New York, second edition, 1991.
- [Ruetsche and Earman, 2005] Laura Ruetsche and John Earman. Relativistic invariance and modal interpretations. Manuscript, University of Pittsburgh, 2005.
- [Ruetsche, 2002] Laura Ruetsche. Interpreting quantum field theory. *Philosophy of Science*, 69(2):348–378, 2002.
- [Ruetsche, 2004] Laura Ruetsche. Intrinsically mixed states: an appreciation. *Studies in the History and Philosophy of Modern Physics*, 35:221–239, 2004.

- [Ryder, 1996] Lewis H. Ryder. *Quantum field theory*. Cambridge University Press, Cambridge, second edition, 1996.
- [Saavedra Rivano, 1972] Neantro Saavedra Rivano. *Catégories Tannakiennes*. Springer-Verlag, Berlin, 1972.
- [Saunders, 1992] Simon Saunders. Locality, complex numbers, and relativistic quantum theory. In David Hull, Micky Forbes, and Kathleen Okruhlik, editors, *PSA; proceedings of the biennial meeting of the Philosophy of Science Association*, volume 1, pages 365–80, East Lansing, MI, 1992. Philosophy of Science Association.
- [Schlieder, 1969] S. Schlieder. Einige Bemerkungen über Projektionsoperatoren (Konsequenzen eines Theorems von Borchers). *Communications in Mathematical Physics*, 13:216–225, 1969.
- [SGA, 1972] *Théorie des topos et cohomologie étale des schémas. Tome 2*. Springer-Verlag, Berlin, 1972. Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par M. Artin, A. Grothendieck et J. L. Verdier. Avec la collaboration de N. Bourbaki, P. Deligne et B. Saint-Donat, Lecture Notes in Mathematics, Vol. 270.
- [Simon, 1996] Barry Simon. *Representations of finite and compact groups*. American Mathematical Society, Providence, RI, 1996.
- [Srinivas, 1980] M. D. Srinivas. Collapse postulate for observables with continuous spectra. *Communications in Mathematical Physics*, 71(2):131–158, 1980.
- [Stinespring, 1955] W. Forrest Stinespring. Positive functions on C^* -algebras. *Proceedings of the American Mathematical Society*, 6:211–216, 1955.
- [Størmer, 1970] Erling Størmer. Asymptotically abelian systems. In *Cargèse lectures in physics, Vol. 4*, pages 195–213. Gordon and Breach, New York, 1970.
- [Streater and Wightman, 1964] R. F. Streater and A. S. Wightman. *PCT, spin and statistics and all that*. W. A. Benjamin, Inc., 1964.
- [Summers and Buchholz, 2005] Stephen J. Summers and Detlev Buchholz. Quantum statistics and locality. *Physics Letters A*, 337:17–21, 2005.
- [Summers and Werner, 1987] Stephen J. Summers and Reinhard Werner. Bell's inequalities and quantum field theory. I. General setting. *Journal of Mathematical Physics*, 28(10):2440–2447, 1987.
- [Summers and Werner, 1988] Stephen J. Summers and Reinhard Werner. Maximal violation of Bell's inequalities for algebras of observables in tangent spacetime regions. *Annales de l'Institut Henri Poincaré. Physique Théorique*, 49(2):215–243, 1988.
- [Summers and Werner, 1995] Stephen J. Summers and Reinhard F. Werner. On Bell's inequalities and algebraic invariants. *Letters in Mathematical Physics*, 33(4):321–334, 1995.
- [Summers, 1990] Stephen J. Summers. On the independence of local algebras in quantum field theory. *Reviews in Mathematical Physics*, 2(2):201–247, 1990.
- [Summers, 1997] Stephen J. Summers. Bell's inequalities and algebraic structure. In *Operator algebras and quantum field theory (Rome, 1996)*, pages 633–646. Internat. Press, Cambridge, MA, 1997.
- [Summers, ND] Stephen Summers. Tomita-Takesaki modular theory. In J. P. Francoise, G. Naber, and T. S. Tsun, editors, *Encyclopedia of Mathematical Physics*. Elsevier, ND.
- [Sunder, 1987] V. S. Sunder. *An invitation to von Neumann algebras*. Springer-Verlag, New York, 1987.
- [Takesaki, 2002] M. Takesaki. *Theory of operator algebras. I*. Springer-Verlag, Berlin, 2002.
- [Takesaki, 2003] M. Takesaki. *Theory of operator algebras. II*. Springer-Verlag, Berlin, 2003.
- [Tannaka, 1939] T. Tannaka. Über den Dualitätssatz der nichtkommutativen topologischen Gruppen. *Tôhoku Mathematics Journal*, 45:1–12, 1939.
- [Teller, 1995] Paul Teller. *An interpretive introduction to quantum field theory*. Princeton University Press, Princeton, NJ, 1995.
- [van Fraassen, 1991] Bas C. van Fraassen. *Quantum mechanics: an empiricist view*. Oxford University Press, New York, 1991.
- [Varadarajan, 2004] V. S. Varadarajan. *Supersymmetry for mathematicians: an introduction*. American Mathematical Society, Providence, RI, 2004.
- [Vermaas, 1999] Pieter E. Vermaas. *A philosopher's understanding of quantum mechanics*. Cambridge University Press, Cambridge, 1999.
- [Wald, 1994] Robert M. Wald. *Quantum field theory in curved spacetime and black hole thermodynamics*. University of Chicago Press, Chicago, IL, 1994.
- [Wallace, forthcoming] David Wallace. In defense of naivete: The conceptual status of Lagrangian QFT. *Synthese*, forthcoming.

- [Wick *et al.*, 1952] G. C. Wick, A. S. Wightman, and E. P. Wigner. The intrinsic parity of elementary particles. *Physical Review*, 88:101–105, 1952.
- [Wightman, 1964] A. S. Wightman. La théorie quantique locale et la théorie quantique des champs. *Ann. Inst. H. Poincaré Sect. A (N.S.)*, 1:403–420, 1964.
- [Wigner, 1939] E. Wigner. On unitary representations of the inhomogeneous Lorentz group. *Ann. of Math. (2)*, 40(1):149–204, 1939.
- [Wizimirski, 1966] Z. Wizimirski. On the existence of the field of operators in the axiomatic quantum field theory. *Bulletin de l'Académie Polonaise des Sciences. Série des Sciences Mathématiques, Astronomiques et Physiques*, 14:91–93, 1966.
- [Wollenberg, 1986] Manfred Wollenberg. Quantum fields as pointlike localized objects. II. *Mathematische Nachrichten*, 128:287–298, 1986.
- [Yngvason, 2005] Jakob Yngvason. The role of type III factors in quantum field theory. *Reports on Mathematical Physics*, 55(1):135–147, 2005.

Appendix.

Abstract Duality Theory for Symmetric Tensor $*$ -Categories

(by Michael Müger)

The aim of this appendix is to give a proof of Theorem 389, first proved by S. Doplicher and J. E. Roberts in 1989, according to which every symmetric tensor $*$ -category with conjugates, direct sums, subobjects and $\text{End } \mathbf{1} = \mathbb{C}$ is equivalent to the category of finite dimensional unitary representations of a uniquely determined compact supergroup. Not much of this material is new, but Theorem 403 probably is, and see Remark 434. However, this seems to be the first exposition of the reconstruction theorem for symmetric tensor categories that gives complete and streamlined proofs, including a short and transparent proof of Tannaka's classical theorem. In the first section we provide the necessary concepts and results of category theory to the extent that they don't involve the notion of fiber functor, whereas the second section is concerned with the Tannaka theory proper. Our main reference for category theory is [Mac Lane, 1998], preferably the second edition. The reader having some experience with categories is advised to skip directly to Section B, using Section A as a reference when needed.

A CATEGORICAL PRELIMINARIES

A.1 Basics

DEFINITION 300. A *category* \mathcal{C} consists of:

- A class $\text{Obj } \mathcal{C}$ of *Objects*. We denote the objects by capital letters X, Y, \dots
- For any two objects X, Y a set $\text{Hom}_{\mathcal{C}}(X, Y)$ of *arrows* (or morphisms); we write $f : X \rightarrow Y$ to indicate that $f \in \text{Hom}_{\mathcal{C}}(X, Y)$, and we omit the subscript \mathcal{C} whenever there is no risk of confusion.
- For any object X a distinguished arrow $\text{id}_X \in \text{End}(X) = \text{Hom}(X, X)$.
- For each $X, Y, Z \in \text{Obj } \mathcal{C}$, a function $\circ : \text{Hom}(Y, Z) \times \text{Hom}(X, Y) \rightarrow \text{Hom}(X, Z)$ such that:

$$h \circ (g \circ f) = (h \circ g) \circ f,$$

and

$$\text{id}_Y \circ f = f, \quad g \circ \text{id}_X = g,$$

whenever $f \in \text{Hom}(X, Y)$, $g \in \text{Hom}(Y, Z)$, and $h \in \text{Hom}(Z, W)$.

DEFINITION 301. A morphism $f \in \text{Hom}(X, Y)$ is an *isomorphism* iff it is invertible, i.e. there is a $g \in \text{Hom}(Y, X)$ such that $g \circ f = \text{id}_X$ and $f \circ g = \text{id}_Y$. If an isomorphism $X \rightarrow Y$ exists, we write $X \cong Y$.

DEFINITION 302. If \mathcal{C} is a category, then a *subcategory* $\mathcal{D} \subset \mathcal{C}$ is defined by a subclass $\text{Obj } \mathcal{D} \subset \text{Obj } \mathcal{C}$ and, for every $X, Y \in \text{Obj } \mathcal{D}$, a subset $\text{Hom}_{\mathcal{D}}(X, Y) \subset \text{Hom}_{\mathcal{C}}(X, Y)$ such that $\text{id}_X \in \text{Hom}_{\mathcal{D}}(X, X)$ for all $X \in \text{Obj } \mathcal{D}$ and the morphisms in \mathcal{D} is closed under the composition \circ of \mathcal{C} . A subcategory $\mathcal{D} \subset \mathcal{C}$ is *full* if $\text{Hom}_{\mathcal{D}}(X, Y) = \text{Hom}_{\mathcal{C}}(X, Y)$ for all $X, Y \in \text{Obj } \mathcal{D}$.

DEFINITION 303. A (covariant) *functor* F from category \mathcal{C} to category \mathcal{D} maps objects of \mathcal{C} to objects of \mathcal{D} and arrows of \mathcal{C} to arrows of \mathcal{D} such that $F(g \circ f) = F(g) \circ F(f)$, and $F(\text{id}_X) = \text{id}_{F(X)}$. A *contravariant functor* is just like a covariant functor except that it reverses the order of arrows.

DEFINITION 304. A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is *faithful*, respectively *full*, if the map

$$F_{X,Y} : \text{Hom}_{\mathcal{C}}(X, Y) \rightarrow \text{Hom}_{\mathcal{D}}(F(X), F(Y))$$

is injective, respectively surjective, for all $X, Y \in \text{Obj } \mathcal{C}$.

DEFINITION 305. A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is *essentially surjective* if for every $Y \in \text{Obj } \mathcal{D}$ there is an $X \in \text{Obj } \mathcal{C}$ such that $F(X) \cong Y$.

DEFINITION 306. If $F : \mathcal{C} \rightarrow \mathcal{D}$ and $G : \mathcal{C} \rightarrow \mathcal{D}$ are functors, then a *natural transformation* η from F to G associates to every $X \in \text{Obj } \mathcal{C}$ a morphism $\eta_X \in \text{Hom}_{\mathcal{D}}(F(X), G(X))$ such that

$$\begin{array}{ccc} F(X) & \xrightarrow{F(s)} & F(Y) \\ \eta_X \downarrow & & \downarrow \eta_Y \\ G(X) & \xrightarrow{G(s)} & G(Y) \end{array}$$

commutes for every arrow $s \in \text{Hom}_{\mathcal{C}}(X, Y)$. If η_X is an isomorphism for every $X \in \text{Obj } \mathcal{C}$, then η is said to be a *natural isomorphism*.

DEFINITION 307. A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is an *equivalence of categories* if there exist a functor $G : \mathcal{D} \rightarrow \mathcal{C}$ and natural isomorphisms $\eta : FG \rightarrow \text{id}_{\mathcal{D}}$ and $\varepsilon : \text{id}_{\mathcal{C}} \rightarrow GF$. Two categories are *equivalent*, denoted $F \simeq G$, if there exists an equivalence $F : \mathcal{C} \rightarrow \mathcal{D}$.

DEFINITION 308. A category is *small* if $\text{Obj } \mathcal{C}$ is a set (rather than just a class). A category is *essentially small* if it is equivalent to a small one, i.e. $\text{Obj } \mathcal{C} / \cong$ is a set.

REMARK 309. Without wanting to go into foundational technicalities we point out that the category of a ‘all representations’ of a group is a huge object. However, considered modulo equivalence the representations are of reasonable cardinality, i.e. are a set.

A.2 Tensor categories and braidings

DEFINITION 310. Given two categories \mathcal{C}, \mathcal{D} , the product category $\mathcal{C} \times \mathcal{D}$ is defined by

$$\begin{aligned} \text{Obj}(\mathcal{C} \times \mathcal{D}) &= \text{Obj}\mathcal{C} \times \text{Obj}\mathcal{D}, \\ \text{Hom}_{\mathcal{C} \times \mathcal{D}}(X \times Y, Z \times W) &= \text{Hom}_{\mathcal{C}}(X, Z) \times \text{Hom}_{\mathcal{D}}(Y, W), \\ \text{id}_{X \times Y} &= \text{id}_X \times \text{id}_Y \end{aligned}$$

with the obvious composition $(a \times b) \circ (c \times d) := (a \circ c) \times (b \circ d)$.

DEFINITION 311. A strict tensor category (or strict monoidal category) is a category \mathcal{C} equipped with a distinguished object $\mathbf{1}$, the tensor unit, and a functor $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ such that:

1. \otimes is associative on objects and morphisms, i.e. $(X \otimes Y) \otimes Z = X \otimes (Y \otimes Z)$ and $(s \otimes t) \otimes u = s \otimes (t \otimes u)$ for all $X, Y, Z, X', Y', Z' \in \text{Obj}\mathcal{C}$ and all $s : X \rightarrow X', t : Y \rightarrow Y', u : Z \rightarrow Z'$.
2. The unit object behaves as it should: $X \otimes \mathbf{1} = X = \mathbf{1} \otimes X$ and $s \otimes \text{id}_{\mathbf{1}} = s = \text{id}_{\mathbf{1}} \otimes s$ for all $s : X \rightarrow Y$.
3. The interchange law

$$(a \otimes b) \circ (c \otimes d) = (a \circ c) \otimes (b \circ d)$$

holds whenever $a \circ c$ and $b \circ d$ are defined.

REMARK 312. Many categories with tensor product are not strict in the above sense. A tensor category is a category equipped with a functor $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$, a unit $\mathbf{1}$ and natural isomorphisms $\alpha_{X,Y,Z} : (X \otimes Y) \otimes Z \rightarrow X \otimes (Y \otimes Z)$, $\lambda_X : \mathbf{1} \otimes X \rightarrow X$, $\rho_X : X \otimes \mathbf{1} \rightarrow X$ satisfying certain identities. The notions of braiding, monoidal functor and monoidal natural transformation generalize to such categories. The generality achieved by considering non-strict categories is only apparent: By the coherence theorems, every (braided/symmetric) tensor category is monoidally naturally equivalent to a strict one. See [Mac Lane, 1998; Joyal and Street, 1993a] for all this.

Strictly speaking (pun intended) the categories of vector spaces and Hilbert spaces are not strict. However, the coherence theorems allow us to pretend that they are, simplifying the formulae considerably. The reader who feels uncomfortable with this is invited to insert the isomorphisms α, λ, ρ wherever they should appear.

DEFINITION 313. A (full) tensor subcategory of a tensor category \mathcal{C} is a (full) subcategory $\mathcal{D} \subset \mathcal{C}$ such that $\text{Obj}\mathcal{D}$ contains the unit object $\mathbf{1}$ and is closed under the tensor product \otimes .

DEFINITION 314. Let \mathcal{C}, \mathcal{D} be strict tensor categories. A tensor functor (or a monoidal functor) is a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ together with isomorphisms $d_{X,Y}^F : F(X) \otimes F(Y) \rightarrow F(X \otimes Y)$ for all $X, Y \in \mathcal{C}$ and a morphism $e^F : \mathbf{1}_{\mathcal{D}} \rightarrow F(\mathbf{1}_{\mathcal{C}})$ such that

1. The morphisms $d_{X,Y}^F$ are natural w.r.t. both arguments.
2. For all $X, Y, Z \in \mathcal{C}$ the following diagram commutes:

$$(56) \quad \begin{array}{ccc} F(X) \otimes F(Y) \otimes F(Z) & \xrightarrow{d_{X,Y}^F \otimes \text{id}_{F(Z)}} & F(X \otimes Y) \otimes F(Z) \\ \downarrow \text{id}_{F(X)} \otimes d_{Y,Z}^F & & \downarrow d_{X \otimes Y, Z}^F \\ F(X) \otimes F(Y \otimes Z) & \xrightarrow{d_{X, Y \otimes Z}^F} & F(X \otimes Y \otimes Z) \end{array}$$

3. The following compositions are the identity morphisms of $F(X)$

$$(57) \quad \begin{array}{l} F(X) \equiv F(X) \otimes \mathbf{1}_{\mathcal{D}} \xrightarrow{\text{id}_{F(X)} \otimes e^F} F(X) \otimes F(\mathbf{1}_{\mathcal{C}}) \xrightarrow{d_{X, \mathbf{1}}} F(X \otimes \mathbf{1}_{\mathcal{C}}) \equiv F(X) \\ F(X) \equiv \mathbf{1}_{\mathcal{D}} \otimes F(X) \xrightarrow{e^F \otimes \text{id}_{F(X)}} F(\mathbf{1}_{\mathcal{C}}) \otimes F(X) \xrightarrow{d_{\mathbf{1}, X}} F(\mathbf{1}_{\mathcal{C}} \otimes X) \equiv F(X) \end{array}$$

for all $X \in \mathcal{C}$.

If \mathcal{C}, \mathcal{D} are tensor $*$ -categories and F is $*$ -preserving, the isomorphisms $e, d_{X,Y}$ are required to be unitary.

DEFINITION 315. Let \mathcal{C}, \mathcal{D} be strict tensor categories and $F, G : \mathcal{C} \rightarrow \mathcal{D}$ tensor functors. A natural transformation $\alpha : \mathcal{C} \rightarrow \mathcal{D}$ is monoidal if

$$\begin{array}{ccc} F(X) \otimes F(Y) & \xrightarrow{d_{X,Y}^F} & F(X \otimes Y) \\ \downarrow \alpha_X \otimes \alpha_Y & & \downarrow \alpha_{X \otimes Y} \\ G(X) \otimes G(Y) & \xrightarrow{d_{X,Y}^G} & G(X \otimes Y) \end{array}$$

commutes for all $X, Y \in \mathcal{C}$ and the composite $\mathbf{1}_{\mathcal{D}} \xrightarrow{e^F} F(\mathbf{1}) \xrightarrow{\alpha_{\mathbf{1}}} G(\mathbf{1})$ coincides with e^G .

REMARK 316. A tensor functor between strict tensor categories is called strict if all the isomorphisms $d_{X,Y}$ and e are identities. However, it is not true that every tensor functor is equivalent to a strict one!

DEFINITION 317. A tensor functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is an *equivalence* (of tensor categories) if there exist a tensor functor $G : \mathcal{D} \rightarrow \mathcal{C}$ and monoidal natural isomorphisms $GF \rightarrow \text{id}_{\mathcal{C}}$ and $FG \rightarrow \text{id}_{\mathcal{D}}$.

PROPOSITION 318. A functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is an equivalence iff F is faithful, full and essentially surjective. A tensor functor $F : \mathcal{C} \rightarrow \mathcal{D}$ of (strict) tensor categories is an equivalence of tensor categories iff F is faithful, full and essentially surjective.

Proof. For the first statement see [Mac Lane, 1998, Theorem 1, p. 91] and for the second [Saavedra Rivano, 1972]. ■

DEFINITION 319. A braiding for a strict tensor category \mathcal{C} is a family of isomorphisms $c_{X,Y} : X \otimes Y \rightarrow Y \otimes X$ for all $X, Y \in \text{Obj } \mathcal{C}$ satisfying

1. Naturality: For every $s : X \rightarrow X', t : Y \rightarrow Y'$, the diagram

$$\begin{array}{ccc}
 X \otimes Y & \xrightarrow{c_{X,Y}} & Y \otimes X \\
 s \otimes t \downarrow & & \downarrow t \otimes s \\
 X' \otimes Y' & \xrightarrow{c_{X',Y'}} & Y' \otimes X'
 \end{array}$$

commutes.

2. The ‘braid equations’ hold, i.e. the diagrams

$$\begin{array}{ccc}
 X \otimes Y \otimes Z & \xrightarrow{c_{X,Y} \otimes \text{id}_Z} & Y \otimes X \otimes Z & \quad & X \otimes Y \otimes Z & \xrightarrow{\text{id}_X \otimes c_{Y,Z}} & X \otimes Z \otimes Y \\
 \searrow c_{X,Y \otimes Z} & & \downarrow \text{id}_Y \otimes c_{X,Z} & & \searrow c_{X \otimes Y, Z} & & \downarrow c_{X,Z} \otimes \text{id}_Y \\
 & & Y \otimes Z \otimes X & & & & Z \otimes X \otimes Y
 \end{array}$$

commute for all $X, Y, Z \in \text{Obj } \mathcal{C}$.

If, in addition, $c_{Y,X} \circ c_{X,Y} = \text{id}_{X \otimes Y}$ holds for all X, Y , the braiding is called a symmetry.

A strict braided (symmetric) tensor category is a strict tensor category equipped with a braiding (symmetry).

DEFINITION 320. If \mathcal{C}, \mathcal{D} are strict braided (symmetric) tensor categories, a tensor functor $F : \mathcal{C} \rightarrow \mathcal{D}$ is braided (symmetric) if


$$F(c_{X,Y}) = c_{F(X),F(Y)} \quad \forall X, Y \in \text{Obj } \mathcal{C}.$$

(Note that on the l.h.s., respectively r.h.s., c is the braiding of \mathcal{C} , respectively \mathcal{D} .)


There is no additional condition for a monoidal natural transformation to be braided/symmetric.

A.3 Graphical notation for tensor categories

We will on some occasions use the so-called ‘tangle diagrams’ for computations in strict (braided) tensor categories, hoping that the reader has seen them before. By way of explanation (for much more detail see e.g. [Kassel, 1995]) we just say that identity morphisms (equivalently, objects) are denoted by vertical lines, a morphism $s : X \rightarrow Y$ by a box with lines corresponding to X and Y entering from below and above, respectively. Compositions and tensor products of morphisms are denoted by vertical and horizontal juxtaposition, respectively. Braiding morphisms are represented by a crossing and the duality morphisms r, \bar{r} by arcs:



(If c is a symmetry, both lines in the braiding are drawn unbroken.) The reason for using this diagrammatic representation is that even relatively simple formulae in tensor categories become utterly unintelligible as soon as morphisms with ‘different numbers of in- and outputs’ are involved, like $s : A \rightarrow B \otimes C \otimes D$. This gets even worse when braidings and duality morphisms are involved. Just one example of a complete formula: The interchange law $s \circ \text{id}_W \circ \text{id}_X \circ t = \text{id}_Y \circ t \circ s \circ \text{id}_Z$ for $s : X \rightarrow Y, t : Z \rightarrow W$ is drawn as



The diagram (correctly!) suggests that we may pull morphisms alongside each other.

A.4 Additive, \mathbb{C} -linear and $*$ -categories

DEFINITION 321. A category is an Ab-category if all hom-sets are abelian groups and the composition \circ is bi-additive.

DEFINITION 322. Let X, Y, Z be objects in a Ab-category. Then Z is called a *direct sum* of X and Y , denoted $Z \cong X \oplus Y$, if there are morphisms $u : X \rightarrow$

$Z, u' : Z \rightarrow X, v : Y \rightarrow Z, v' : Z \rightarrow Y$ such that $u' \circ u = \text{id}_X, v' \circ v = \text{id}_Y$ and $u \circ u' + v \circ v' = \text{id}_Z$. (Note that every $Z' \cong Z$ also is a direct sum of X and Y . Thus direct sums are defined only up to isomorphism, which is why we don't write $Z = X \oplus Y$.) We say that \mathcal{C} has *direct sums* if there exists a direct sum $Z \cong X \oplus Y$ for any two object X, Y .

DEFINITION 323. An object $\mathbf{0}$ in a category \mathcal{C} is called a zero object if, for every $X \in \mathcal{C}$, the sets $\text{Hom}(X, \mathbf{0})$ and $\text{Hom}(\mathbf{0}, X)$ both contain precisely one element. A morphism to or from a zero object is called a zero morphism.

It is immediate from the definition that any two zero objects are isomorphic. If a category doesn't have a zero object it is straightforward to add one. If z is a zero morphism and f is any morphism, then $z \circ f, f \circ z, z \otimes f, f \otimes z$ are zero morphisms (provided they make sense).

DEFINITION 324. An *additive* category is an Ab-category that has a zero object and direct sums.

EXAMPLE 325. The category of abelian groups (with the trivial group as zero).

DEFINITION 326. A category \mathcal{C} is called \mathbb{C} -linear if $\text{Hom}(X, Y)$ is a \mathbb{C} -vector space for all $X, Y \in \text{Obj } \mathcal{C}$ and the composition map $\circ : (f, g) \mapsto g \circ f$ is bilinear. If \mathcal{C} is a tensor category we require that also $\otimes : (f, g) \mapsto g \otimes f$ is bilinear. Functors between \mathbb{C} -linear category are always assumed to be \mathbb{C} -linear, i.e. $\text{Hom}_{\mathcal{C}}(X, Y) \rightarrow \text{Hom}_{\mathcal{D}}(F(X), F(Y))$ must be \mathbb{C} -linear.

DEFINITION 327. A *positive *-operation* on a \mathbb{C} -linear category is a family of maps that to every morphism $s \in \text{Hom}(X, Y)$ associates a morphism $s^* \in \text{Hom}(Y, X)$. This map must be antilinear, involutive ($(s^*)^* = s$) and positive in the sense that $s^* \circ s = 0$ implies $s = 0$. A **-category* is a \mathbb{C} -linear category equipped with a positive *-operation. A *tensor *-category* is a tensor category with a positive *-operation such that $(s \otimes t)^* = s^* \otimes t^*$ for all s, t . We consider only unitary braidings (symmetries) of tensor *-categories.

DEFINITION 328. A morphism $v : X \rightarrow Y$ in a *-category is called an *isometry* if $v^* \circ v = \text{id}_X$. An isometry v is called a *unitary* if it satisfies $v \circ v^* = \text{id}_Y$. A morphism $p \in \text{End } X$ is called a projector if $p = p \circ p = p^*$. We say that \mathcal{C} has *subobjects* if for every projector $p \in \text{End } X$ there exists an isometry $v : Y \rightarrow X$ such that $v \circ v^* = p$. In a *-category we strengthen Definition 322 by requiring that $u' = u^*, v' = v^*$, i.e. u, v must be isometries.

DEFINITION 329. A functor F between *-categories is **-preserving* if $F(s^*) = F(s)^*$ for every morphism s . The isomorphisms $d_{X,Y}, e$ coming with a functor between tensor *-categories coming with a functor of tensor *-categories are required to be unitaries.

DEFINITION 330. Let \mathcal{C} be a tensor *-category and $X \in \text{Obj } \mathcal{C}$. An object $\overline{X} \in \text{Obj } \mathcal{C}$ is called a *conjugate* object of X if there exist morphisms $r : \mathbf{1} \rightarrow \overline{X} \otimes X$

and $r : 1 \rightarrow X \otimes \bar{X}$ satisfying the ‘conjugate equations’

$$\begin{aligned} \text{id}_X \otimes r^* \circ \bar{r} \otimes \text{id}_X &= \text{id}_X, \\ \text{id}_{\bar{X}} \otimes \bar{r}^* \circ r \otimes \text{id}_{\bar{X}} &= \text{id}_{\bar{X}}. \end{aligned}$$

We say that (\bar{X}, r, \bar{r}) is a *conjugate* of X . If every non-zero object of \mathcal{C} has a conjugate then we say that \mathcal{C} has conjugates.

Note also that a zero object cannot have a conjugate. If $(\bar{X}, r, \bar{r}), (\bar{X}', r', \bar{r}')$ both are conjugates of X then one easily verifies that $\text{id}_{\bar{X}'} \otimes \bar{r}^* \circ r' \otimes \text{id}_{\bar{X}} : \bar{X} \rightarrow \bar{X}'$ is unitary. Thus the conjugate is unique up to unitary equivalence.

DEFINITION 331. An object X in a \mathbb{C} -linear category is *irreducible* if $\text{End } X = \text{Cid}_X$.

DEFINITION 332. A TC^* is a tensor $*$ -category with finite dimensional hom-sets, with conjugates, direct sums, subobjects and irreducible unit 1 . A BTC^* is a TC^* with a unitary braiding. An STC^* is a TC^* with a unitary symmetry.

EXAMPLE 333. The tensor $*$ -category \mathcal{H} of finite dimensional Hilbert spaces is a STC^* . The symmetry $c_{H,H'} : H \otimes H' \rightarrow H' \otimes H$ is given by the flip isomorphism $\Sigma : x \otimes y \mapsto y \otimes x$. The conjugate of an object H is the Hilbert space dual \bar{H} . Picking a basis $\{e_i\}$ of H with dual basis $\{f_i\}$, the conjugation morphisms are given by

$$r = \sum_i f_i \otimes e_i, \quad \bar{r} = \sum_i e_i \otimes f_i.$$

In the same way one sees that the category $\text{Rep}_f G$ of finite dimensional unitary representations of a compact group G is an STC^* .

LEMMA 334. A TC^* is semisimple, i.e. every object is a finite direct sum of irreducible objects.

Proof. For every $X \in \mathcal{C}$, $\text{End } X$ is a finite dimensional \mathbb{C} -algebra with a positive involution. Such an algebra is semisimple, to wit a multi matrix algebra. Thus id_X is a sum of projections p_i that are minimal in the sense that $p_i \text{End } X p_i \cong \mathbb{C}$. Since \mathcal{C} has subobjects, there are objects X_i corresponding to the p_i , which are irreducible by minimality of the p_i . Clearly, $X \cong \oplus_i X_i$. ■

DEFINITION 335. A solution (X, r, \bar{r}) of the conjugate equations is called standard if

$$r^* \circ \text{id}_{\bar{X}} \otimes s \circ r = \bar{r}^* \circ s \otimes \text{id}_{\bar{X}} \circ \bar{r}$$

for all $s \in \text{End } X$. In this case, (\bar{X}, r, \bar{r}) is called a standard conjugate.

LEMMA 336. Let \mathcal{C} be a TC^* and (\bar{X}, r, \bar{r}) a conjugate for $X \in \mathcal{C}$. Let $v_i : X_i \rightarrow X, w_i : \bar{X}_i \rightarrow \bar{X}$ be isometries effecting the direct sum decomposition of X, \bar{X} into irreducibles. Then (\bar{X}, r, \bar{r}) is a standard conjugate iff $(\bar{X}_i, w_i^* \otimes v_i^* \circ r, v_i^* \otimes w_i^* \circ \bar{r})$ is a standard conjugate for X_i for all i . Every object admits a standard conjugate.

Proof. For the equivalence claim, see [Longo and Roberts, 1997], in particular Lemma 3.9. (Note that in [Longo and Roberts, 1997], standardness is defined by the property in the statement above.) We use this to prove that every objects admits a standard conjugate. If X is irreducible, we have $\text{End } X = \mathbb{C}\text{id}_X$. Therefore the standardness condition reduces to $r^* \circ r = \bar{r}^* \circ \bar{r}$, thus a conjugate (\bar{X}, r, \bar{r}) can be made standard by rescaling r, \bar{r} . In the general case, we use semisimplicity to find a direct sum decomposition of X into irreducibles X_i . Let $(\bar{X}_i, r_i, \bar{r}_i)$ be standard conjugates of the X_i and put $\bar{X} = \bigoplus \bar{X}_i$. Let $v_i : X_i \rightarrow X$, $w_i : \bar{X}_i \rightarrow \bar{X}$ be the isometries effecting the direct sums. Defining $r = \sum_i w_i \otimes v_i \circ r_i$ and $\bar{r} = \sum_i v_i \otimes w_i \circ \bar{r}_i$, the criterion in the first part of the lemma applies and gives standardness of (\bar{X}, r, \bar{r}) . ■

LEMMA 337. *Let (\bar{X}, r, \bar{r}) be a (standard) conjugate of X , let $p \in \text{End } X$ a projection and define $\bar{p} = r^* \otimes \text{id}_{\bar{X}} \circ \text{id}_{\bar{X}} \otimes p \otimes \text{id}_{\bar{X}} \circ \text{id}_{\bar{X}} \otimes \bar{r} \in \text{End } \bar{X}$. If $v : Y \rightarrow X$, $w : \bar{Y} \rightarrow \bar{X}$ are isometries such that $v \circ v^* = p, w \circ w^* = \bar{p}$ then $(\bar{Y}, w^* \otimes v^* \circ r, v^* \otimes w^* \circ \bar{r})$ is a (standard) conjugate for Y .*

Proof. Omitted. For the easy proof see [Longo and Roberts, 1997] or [Müger, 2000]. ■

LEMMA 338. *If $(\bar{X}, r, \bar{r}), (\bar{Y}, r', \bar{r}')$ are (standard) conjugates of X, Y , respectively, then $(\bar{Y} \otimes \bar{X}, r'', \bar{r}'')$, where $r'' = \text{id}_{\bar{Y}} \otimes r \otimes \text{id}_Y \circ r', \bar{r}'' = \text{id}_X \otimes \bar{r}' \otimes \text{id}_{\bar{X}} \circ \bar{r}$ is a (standard) conjugate for $X \otimes Y$.*

Proof. That $(\bar{Y} \otimes \bar{X}, r'', \bar{r}'')$ is a conjugate is an easy computation. Standardness is less obvious since the map $\text{End } X \otimes \text{End } Y \rightarrow \text{End } X \otimes Y$ need not be surjective. However, it follows using the alternative characterization of standardness given in Lemma 336. ■

PROPOSITION 339. *Let \mathcal{C} be a TC^* . Let $X \in \mathcal{C}$ and let (\bar{X}, r, \bar{r}) be a standard conjugate. Then the map*

$$\text{Tr}_X : \text{End } X \rightarrow \mathbb{C}, \quad s \mapsto r^* \circ \text{id}_{\bar{X}} \otimes s \circ r$$

is well defined, i.e. independent of the choice of (\bar{X}, r, \bar{r}) . It is called the trace. It satisfies

$$\begin{aligned} \text{Tr}_X(s \circ t) &= \text{Tr}_Y(t \circ s) \quad \forall s : Y \rightarrow X, t : X \rightarrow Y, \\ \text{Tr}_{X \otimes Y}(s \otimes t) &= \text{Tr}_X(s) \text{Tr}_Y(t) \quad \forall s \in \text{End } X, t \in \text{End } Y. \end{aligned}$$

Proof. Easy exercise. ■

DEFINITION 340. Let \mathcal{C} be a TC^* and $X \in \mathcal{C}$. The *dimension* of X is defined by $d(X) = \text{Tr}_X(\text{id}_X)$, i.e. $d(X) = r^* \circ r$ for any standard conjugate (\bar{X}, r, \bar{r}) .

LEMMA 341. *The dimension is additive ($d(X \oplus Y) = d(X) + d(Y)$) and multiplicative ($d(X \otimes Y) = d(X)d(Y)$). Furthermore, $d(\bar{X}) = d(X) \geq 1$ for every object, and $d(X) = 1$ implies that $X \otimes \bar{X} \cong \mathbf{1}$, i.e. X is invertible.*

Proof. Additivity is immediate by the discussion of standard conjugates. Multiplicativity of the dimension follows from Lemma 338.

If $(\overline{X}, r, \bar{r})$ is a standard conjugate for X , then (X, \bar{r}, r) is a standard conjugate for \overline{X} , implying $d(\overline{X}) = d(X)$. The positivity of the $*$ -operation implies that $d(X) = r^* \circ r > 0$. Since $X \otimes \overline{X}$ contains $\mathbf{1}$ as a direct summand, we have $d(X)^2 \geq 1$, thus $d(X) \geq 1$. Finally, if $d(X) = 1$, $\mathbf{1}$ is the only direct summand of $X \otimes \overline{X}$, to wit $X \otimes \overline{X} \cong \mathbf{1}$. Similarly, $\overline{X} \otimes X \cong \mathbf{1}$. ■

DEFINITION 342. Let \mathcal{C} be a BTC^* and $X \in \mathcal{C}$. The twist $\Theta(X) \in \text{End } X$ is defined by

$$\Theta(X) = r^* \otimes \text{id}_X \circ \text{id}_{\overline{X}} \otimes c_{X,X} \circ r \otimes \text{id}_X,$$

where $(\overline{X}, r, \bar{r})$ is a standard solution of the conjugate equations.

LEMMA 343. Let \mathcal{C} be a BTC^* . Then

- (i) $\Theta(X)$ is well defined, i.e. does not depend on the choice of $(\overline{X}, r, \bar{r})$.
- (ii) For every morphism $s : X \rightarrow Y$ we have $\Theta(Y) \circ s = s \circ \Theta(X)$. (I.e., Θ is a natural transformation of the identity functor of \mathcal{C} .)
- (iii) $\Theta(X)$ is unitary.
- (iv) $\Theta(X \otimes Y) = \Theta(X) \otimes \Theta(Y) \circ c_{Y,X} \circ c_{X,Y}$ for all X, Y .
- (v) If \mathcal{C} is an STC^* , this simplifies to $\Theta(X)^2 = \text{id}_X$ and $\Theta(X \otimes Y) = \Theta(X) \otimes \Theta(Y)$ for all $X, Y \in \mathcal{C}$ (i.e., Θ is a monoidal natural transformation of the identity functor of \mathcal{C}). If X, Y are irreducible, we have $\omega(X) = \pm 1$ and $\omega_Z = \omega_X \omega_Y$ for all irreducible direct summands $Z \prec X \otimes Y$.


Proof. (i) is proven as Proposition 339. The other verifications are not-too-difficult computations, for which we refer to [Longo and Roberts, 1997] or [Müger, 2000]. We just comment on (v): In an STC^* we have $c_{X,X}^* = c_{X,X}^{-1} = c_{X,X}$, implying $\Theta(X)^* = \Theta(X)$. Together with unitarity this gives $\Theta(X)^2 = \text{id}_X$. Multiplicativity of Θ in an STC^* follows from $c_{Y,X} \circ c_{X,Y} = \text{id}$. If X, Y are irreducible, we have $\Theta(X) = \omega_X \text{id}_X$, $\Theta(Y) = \omega_Y \text{id}_Y$ and thus $\Theta(X \otimes Y) = \omega_X \omega_Y \text{id}_{X \otimes Y}$. Now $\omega(Z) = \omega_X \omega_Y$ for irreducible $Z \prec X \otimes Y$ follows by naturality of Θ . ■

The following is a reworking of Propositions 4.4 and 4.5 in [Longo and Roberts, 1997].

PROPOSITION 344. Let \mathcal{C}, \mathcal{D} be BTC^* s and $E : \mathcal{C} \rightarrow \mathcal{D}$ a $*$ -preserving braided tensor functor. If $(\overline{X}, r, \bar{r})$ is a standard conjugate of $X \in \mathcal{C}$, then $(E(\overline{X}), (d_{\overline{X}, X}^E)^{-1} \circ E(r) \circ e^E, (d_{X, \overline{X}}^E)^{-1} \circ E(\bar{r}) \circ e^E)$ is a standard conjugate for $E(X)$. In particular,

$$d(E(X)) = d(X), \quad \Theta(E(X)) = E(\Theta(X)) \quad \forall X \in \mathcal{C}.$$

Proof. We assume for a while that the functor E is strict and that X is irreducible. Let (\bar{X}, r, \bar{r}) be a standard conjugate. Since E preserves the conjugate equations, $(E(\bar{X}), E(r), E(\bar{r}))$ is a conjugate for $E(X)$, but if E is not full, standardness requires proof. We begin with



Thus $c_{\bar{X}, X}^* \circ \bar{r} = \omega_{\bar{X}} \cdot r$, which is equivalent to $c_{\bar{X}, X} \circ r = \bar{\omega}_{\bar{X}} \bar{r}$. Now we let $s \in \text{End } E(X)$ and compute

$$\begin{aligned}
 E(r^*) \circ \text{id}_{E(\bar{X})} \otimes s \circ E(r) &= \\
 &= E(r^*) \circ c_{E(\bar{X}), E(X)}^* \circ c_{E(\bar{X}), E(X)} \circ \text{id}_{E(\bar{X})} \otimes s \circ E(r) \\
 &= (c_{E(\bar{X}), E(X)} \circ E(r))^* \circ c_{E(\bar{X}), E(X)} \circ \text{id}_{E(\bar{X})} \otimes s \circ E(r) \\
 &= (c_{E(\bar{X}), E(X)} \circ E(r))^* \circ s \otimes \text{id}_{E(\bar{X})} \circ c_{E(\bar{X}), E(X)} \circ E(r) \\
 &= E(c_{\bar{X}, X} \circ r)^* \circ s \otimes \text{id}_{E(\bar{X})} \circ E(c_{\bar{X}, X} \circ r) \\
 &= E(\bar{\omega}_{\bar{X}} \bar{r})^* \circ s \otimes \text{id}_{E(\bar{X})} \circ E(\bar{\omega}_{\bar{X}} \bar{r}) \\
 &= E(\bar{r})^* \circ s \otimes \text{id}_{E(\bar{X})} \circ E(\bar{r}),
 \end{aligned}$$

which means that $(E(\bar{X}), E(r), E(\bar{r}))$ is a standard conjugate for $E(X)$. (We have used unitarity of the braiding, the fact that E is $*$ -preserving and braided, $c_{\bar{X}, X} \circ r = \bar{\omega}_{\bar{X}} \bar{r}$ and $|\omega_{\bar{X}}| = 1$.)

Now let X be reducible, (\bar{X}, r, \bar{r}) a standard conjugate and let $v_i : X_i \rightarrow X$, $w_i : \bar{X}_i \rightarrow \bar{X}$ be isometries effecting the decompositions into irreducibles. Defining $r_i = w_i^* \otimes v_i^* \circ r$, $\bar{r}_i = v_i^* \otimes w_i^* \circ \bar{r}$, $(\bar{X}_i, r_i, \bar{r}_i)$ is standard by Lemma 336. Thus $(E(\bar{X}_i), E(r_i), E(\bar{r}_i))$ is standard by the first half of this proof. In view of $E(r) = E(\sum_i w_i \otimes v_i \circ r_i) = \sum_i E(w_i \otimes v_i) \circ E(r_i)$ and similarly for $E(\bar{r})$, it follows that $(E(\bar{X}), E(r), E(\bar{r}))$ is standard (since it is a direct sum of standard conjugates).

If E is not strict, we have to insert the unitaries $d_{X, Y}^E : E(X) \otimes E(Y) \rightarrow E(X \otimes Y)$, $e^E : \mathbf{1} \rightarrow E(\mathbf{1})$ at the obvious places in the above computations, but nothing else changes. That E preserves dimensions follows since the dimension is defined in terms of a standard conjugate. Finally, standardness of $(E(\bar{X}), E(r), E(\bar{r}))$ together with $E(c_{X, Y}) = c_{E(X), E(Y)}$ imply $\Theta(E(X)) = E(\Theta(X))$. ■

We close this subsection by commenting on the relation of $*$ -categories with the more general notion of C^* -tensor categories of [Doplicher and Roberts, 1989; Longo and Roberts, 1997].

DEFINITION 345. A C^* -category is a \mathbb{C} -linear category with a positive $*$ -operation, where $\text{Hom}(X, Y)$ is a Banach space for all X, Y and $\|s \circ t\|_{\text{Hom}(X, Z)} \leq \|s\|_{\text{Hom}(X, Y)} \cdot \|t\|_{\text{Hom}(Y, Z)}$ for all $s : X \rightarrow Y, t : Y \rightarrow Z$ and $\|s^* \circ s\|_{\text{End } X} = \|s\|_{\text{Hom}(X, Y)}^2$ for all $s : X \rightarrow Y$. (Thus each $\text{End } X$ is a C^* -algebra.) A C^* -tensor category is a C^* -category and a tensor category such that $\|s \otimes t\| \leq \|s\| \cdot \|t\|$ for all s, t .

PROPOSITION 346. [Longo and Roberts, 1997] *Let \mathcal{C} be a C^* -tensor category with direct sums and irreducible unit. Whenever $X, Y \in \mathcal{C}$ admit conjugates then $\dim \text{Hom}(X, Y) < \infty$. Thus a C^* -tensor category with direct sums, subobjects, conjugates and irreducible unit is a TC^* . Conversely, given a TC^* , there are unique norms on the spaces $\text{Hom}(X, Y)$ rendering \mathcal{C} a C^* -tensor category.*

Proof. Assume that $X \in \mathcal{C}$ has a conjugate (X, r, \bar{r}) . Then the map $\text{End } X \rightarrow \text{Hom}(\mathbf{1}, \bar{X} \otimes X), s \mapsto \text{id}_{\bar{X}} \otimes s \circ r$ is an isomorphism of vector spaces since $t \mapsto \bar{r}^* \circ \text{id}_X \circ \text{id}_X \otimes t$ is its inverse, as is verified using the conjugate equations. Now, $\text{Hom}(\mathbf{1}, \bar{X} \otimes X)$ is a pre-Hilbert space w.r.t. the inner product $\langle a, b \rangle_{\text{id}_1} = a^* \circ b$, and it is complete because \mathcal{C} is a C^* -tensor category. Choose an orthogonal basis $(e_i)_{i \in I}$ in $\text{Hom}(\mathbf{1}, \bar{X} \otimes X)$. Then each $e_i : \mathbf{1} \rightarrow \bar{X} \otimes X$ is an isometry and $e_i^* \circ e_j = 0$ for $i \neq j$, implying that $\bar{X} \otimes X$ contains $\#I$ copies of $\mathbf{1}$ as direct summands. Since X has a conjugate, so does $\bar{X} \otimes X$, but this is impossible if $\#I$ is infinite. Thus $\text{Hom}(\mathbf{1}, \bar{X} \otimes X)$ and therefore $\text{End } X$ is finite dimensional.

Given arbitrary X, Y having conjugates, pick a direct sum $Z \cong X \oplus Y$ with isometries $u : X \rightarrow Z, v : Y \rightarrow Z$. Then also Z has a conjugate, cf. Lemma 336, and therefore $\dim \text{End } Z < \infty$. Now, the map $\text{Hom}(X, Y) \rightarrow \text{End } Z$ given by $s \mapsto v \circ s \circ u^*$ is injective since it has $t \mapsto v^* \circ t \circ u$ as inverse. This implies $\dim \text{Hom}(X, Y) < \infty$.

We omit the proof of the implication $TC^* \Rightarrow C^*$ -tensor category, since it will not be used in the sequel. It can be found in [Müger, 2000]. ■

This result shows that the assumptions made in Appendix B are equivalent to those of [Doplicher and Roberts, 1989], formulated in terms of C^* -tensor categories.

A.5 Abelian categories

In the second half of Appendix B, which is of a purely algebraic nature, we will need some basic facts from the theory of abelian categories. Good references are, e.g., [Gabriel, 1962] and [Mac Lane, 1998, Chapter VIII].

DEFINITION 347. A morphism $s : X \rightarrow Y$ is called monic if $s \circ t_1 = s \circ t_2$ implies $t_1 = t_2$, whenever t_1, t_2 are morphisms with target X and the same source. A morphism $s : X \rightarrow Y$ is called epi if $t_1 \circ s = t_2 \circ s$ implies $t_1 = t_2$, whenever t_1, t_2 are morphisms with source Y and the same target.

DEFINITION 348. Let \mathcal{C} be an additive category. Given a morphism $f : X \rightarrow Y$, a morphism $k : Z \rightarrow X$ is a kernel of f if $f \circ k = 0$ and given any morphism $k' : Z' \rightarrow X$ such that $f \circ k' = 0$, there is a unique morphism $l : Z' \rightarrow Z$ such that $k' = k \circ l$.

A cokernel of $f : X \rightarrow Y$ is a morphism $c : Y \rightarrow Z$ if $c \circ f = 0$ and given any morphism $c' : Y \rightarrow Z'$ such that $c' \circ f = 0$, there is a unique $d : Z \rightarrow Z'$ such that $c' = d \circ c$.

It is an easy consequence of the definition that every kernel is monic and every cokernel is epic.

DEFINITION 349. An additive category \mathcal{C} is *abelian* if

1. Every morphism has a kernel and a cokernel.
2. Every monic morphism is the kernel of some morphism.
3. Every epic morphism is the cokernel of some morphism.

PROPOSITION 350. Let \mathcal{C} be an abelian category. Then

- (i) Every monic is the kernel of its cokernel and every epi is the cokernel of its kernel.
- (ii) A morphism is an isomorphism iff it is monic and epic.
- (iii) Every morphism $f : X \rightarrow Y$ in an abelian category admits a factorization $f = m \circ e$, where $e : X \rightarrow Z$ is epi and $m : Z \rightarrow Y$ is monic. Given another epi $e' : X \rightarrow Z'$ and monic $m' : Z' \rightarrow Y$ such that $f = m' \circ e'$, there exists an isomorphism $u : Z \rightarrow Z'$ such that $e' = u \circ e$ and $m = m' \circ u$.

Proof. See [Mac Lane, 1998, Chapter VIII] for detailed proofs. The ‘only if’ of (ii) is trivial. Concerning (iii): Defining $m = \ker(\text{coker}(f))$, m is monic. In view of $(\text{coker } f) \circ f = 0$, f factors as $f = m \circ e$ for a unique e . Next one proves that e is epi and $e = \text{coker}(\ker(f))$. ■

DEFINITION 351. The *image* of a morphism $f : X \rightarrow Y$ in an abelian category is the monic $m : Z \rightarrow Y$ (unique up to isomorphism) in the monic-epic factorization $X \xrightarrow{e} Z \xrightarrow{m} Y$ of f .

In a concrete abelian category, the object Z is isomorphic to the usual image of f , which is a subset of Y , whence the terminology.

DEFINITION 352. An object P in an abelian category is *projective* if, given any epimorphism $p : A \rightarrow B$ and any morphism $b : P \rightarrow B$ there is a morphism $a : P \rightarrow A$ such that $b = p \circ a$.

LEMMA 353. Any TC^* \mathcal{C} that has a zero object is abelian.

Proof. It is clear that \mathcal{C} is additive. The other requirements of Definition 349 follow with a little work from semisimplicity, cf. Lemma 334. ■

A.6 Commutative algebra in abelian symmetric tensor categories

A considerable part of the well known algebra of commutative rings, their ideals and modules (living in the category Ab of abelian groups) can be generalized to other abelian symmetric or even braided tensor categories. We state just those facts that will be needed, some of which might be new.

DEFINITION 354. Let \mathcal{D} be a strict tensor category. Then a monoid in \mathcal{D} is a triple (Q, m, η) , where $Q \in \mathcal{D}$ and $m : Q \otimes Q \rightarrow Q$ and $\eta : \mathbf{1} \rightarrow Q$ are morphisms satisfying

$$m \circ (m \otimes \text{id}_Q) = m \circ (\text{id}_Q \otimes m), \quad m \circ \eta \otimes \text{id}_Q = \text{id}_Q = m \circ \text{id}_Q \otimes \eta.$$

If \mathcal{D} is braided then the monoid is called commutative if $m \circ c_{Q,Q} = m$.

DEFINITION 355. Let (Q, m, η) be a monoid in the strict tensor category \mathcal{D} . Then a Q -module (in \mathcal{D}) is a pair (M, μ) , where $M \in \mathcal{D}$ and $\mu : Q \otimes M \rightarrow M$ satisfy

$$\mu \circ \text{id}_Q \otimes \mu = \mu \circ m \otimes \text{id}_M, \quad \mu \circ \eta \otimes \text{id}_M = \text{id}_M.$$

A morphism $s : (M, \mu) \rightarrow (R, \rho)$ of Q -modules is a morphism $s \in \text{Hom}_{\mathcal{D}}(M, R)$ satisfying $s \circ \mu = \rho \circ \text{id}_Q \otimes s$. The Q -modules in \mathcal{D} and their morphisms form a category $Q - \text{Mod}_{\mathcal{D}}$. If \mathcal{D} is k -linear then $Q - \text{Mod}_{\mathcal{D}}$ is k -linear. The hom-sets in the category $Q - \text{Mod}$ are denoted by $\text{Hom}_Q(\cdot, \cdot)$.

REMARK 356.

1. The preceding definitions, which are obvious generalizations of the corresponding notions in Vect , generalize in a straightforward way to non-strict tensor categories.
2. If (M, μ) is a Q -module and $X \in \mathcal{D}$ then $(Q \otimes X, \mu \otimes \text{id}_X)$ is a Q -module.
3. If \mathcal{D} has direct sums, we can define the direct sum (R, ρ) of two Q -modules $(M_1, \mu_1), (M_2, \mu_2)$. Concretely, if $v_i : M_i \rightarrow R, i = 1, 2$ are the isometries corresponding to $R \cong M_1 \oplus M_2$ then $\rho = v_1 \circ \mu_1 \circ \text{id}_Q \otimes v_1^* + v_2 \circ \mu_2 \circ \text{id}_Q \otimes v_2^*$ provides a Q -module structure.
4. Given a monoid (Q, m, η) in \mathcal{D} , we have an obvious Q -module (Q, m) , and for any $n \in \mathbb{N}$ we can consider $n \cdot (Q, m)$, the direct sum of n copies of the Q -module (Q, m) .

DEFINITION 357. Let \mathcal{D} be a strict tensor category with unit $\mathbf{1}$ and let (Q, m, η) be a monoid in \mathcal{D} . We define a monoid Γ_Q in the category of sets by $\Gamma_Q = \text{Hom}(\mathbf{1}, Q)$, the multiplication being given by $s \bullet t = m \circ t \otimes s$ and the unit by η . If \mathcal{D} is braided and (Q, m, η) commutative then Γ_Q is commutative.

LEMMA 358. *Let \mathcal{D} be a strict tensor category and (Q, m, η) a monoid in \mathcal{D} . Then there is an isomorphism of monoids $\gamma : \text{End}_Q((Q, m)) \rightarrow (\Gamma_Q, \bullet, \eta)$ given by*

$$\begin{aligned} \gamma : \text{End}_Q((Q, m)) &\rightarrow \text{Hom}(\mathbf{1}, Q), & u &\mapsto u \circ \eta, \\ \gamma^{-1} : \text{Hom}(\mathbf{1}, Q) &\rightarrow \text{End}_Q((Q, m)), & s &\mapsto m \circ \text{id}_Q \otimes s. \end{aligned}$$

If \mathcal{D} (and thus $Q - \text{Mod}_{\mathcal{D}}$) is k -linear then γ is an isomorphism of k -algebras. If \mathcal{D} is braided and the monoid (Q, m, η) is commutative then the monoid (k -algebra) $(\Gamma_Q, \bullet, \eta)$, and therefore also $\text{End}_Q((Q, m))$, is commutative.

Proof. That $(\Gamma_Q, \bullet, \eta)$ is a monoid (resp. associative k -algebra) is immediate since (Q, m, η) is a monoid. For $s \in \text{Hom}(\mathbf{1}, Q)$ we have $\gamma(\gamma^{-1}(s)) = m \circ \text{id}_Q \otimes s \circ \eta = s$ by the monoid axioms. On the other hand, for $u \in \text{End}_Q((Q, m))$ we have

$$\gamma^{-1}(\gamma(u)) = m \circ \text{id}_Q \otimes (u \circ \eta) = m \circ \text{id}_Q \otimes u \circ \text{id}_Q \otimes \eta = u \circ m \circ \text{id}_Q \otimes \eta = u,$$

where the third equality is due to the fact that s is a Q -module map (cf. Definition 355). Clearly $\gamma(\text{id}_Q) = \eta$. Furthermore,

$$\begin{aligned} \gamma^{-1}(s) \circ \gamma^{-1}(t) &= (m \circ \text{id}_Q \otimes s) \circ (m \circ \text{id}_Q \otimes t) = m \circ m \otimes \text{id}_Q \circ \text{id}_Q \otimes t \otimes s \\ &= m \circ \text{id}_Q \otimes m \circ \text{id}_Q \otimes t \otimes s = \gamma^{-1}(s \bullet t). \end{aligned}$$

If \mathcal{D} is braided and the monoid (Q, m, η) is commutative then

$$s \bullet t = m \circ t \otimes s = m \circ c_{Q, Q} \circ s \otimes t = m \circ s \otimes t = t \bullet s,$$

where we used naturality of the braiding and commutativity of the monoid. ■

REMARK 359. 1. We have seen that a monoid (Q, m, η) in any abstract tensor category gives rise to a monoid $(\Gamma_Q, \bullet, \eta)$ that is concrete, i.e. lives in the category Sets. The latter has the cartesian product as a tensor product and any one-element set is a tensor unit $\mathbf{1}$. Thus for any $X \in \text{Sets}$, $\text{Hom}(\mathbf{1}, X)$ is in bijective correspondence to the elements of X . Therefore, if $\mathcal{D} = \text{Sets}$ then the monoids (Q, m, η) and $(\Gamma_Q, \bullet, \eta)$ are isomorphic. For this reason, we call Γ_Q the monoid of elements of Q even when \mathcal{D} is an abstract category.

2. The commutativity of $\text{End}_Q((Q, m))$ in the case of a commutative monoid (Q, m, η) in a braided tensor category \mathcal{D} has a very natural interpretation: If \mathcal{D} has coequalizers, which holds in any abelian category, then the category $Q - \text{Mod}_{\mathcal{D}}$ is again a tensor category and the Q -module (Q, m) is its unit object. In any tensor category with unit $\mathbf{1}$, $\text{End } \mathbf{1}$ is a commutative monoid (commutative k -algebra if \mathcal{D} is k -linear). This is the real reason why $\text{End}_Q((Q, m))$ is commutative. More is known: If \mathcal{D} is symmetric and Q abelian, then the tensor category $Q - \text{Mod}_{\mathcal{D}}$ is again symmetric. (In the braided case this need not be true, but $Q - \text{Mod}_{\mathcal{D}}$ always has a distinguished full subcategory that is braided.)

We now specialize to abelian categories.

PROPOSITION 360. *Let (Q, m, η) be a monoid in an abelian strict tensor category \mathcal{D} . Then the category $Q - \text{Mod}_{\mathcal{D}}$ is abelian.*

Proof. Omitted. (This is a nice exercise on abelian categories.) ■

DEFINITION 361. Let \mathcal{D} be an abelian strict symmetric tensor category. An ideal in a commutative monoid (Q, m, η) is a monic $j : (J, \mu_J) \rightarrow (Q, m)$ in the category $Q - \text{Mod}$. An ideal $j : (J, \mu_J) \rightarrow (Q, m)$ is called *proper* if j is not an isomorphism (i.e. not epi). If $j : (J, \mu_J) \rightarrow (Q, m)$ and $j' : (J', \mu_{J'}) \rightarrow (Q, m)$ are ideals then $j : (J, \mu_J) \rightarrow (Q, m)$ is *contained* in $j' : (J', \mu_{J'}) \rightarrow (Q, m)$, denoted $j < j'$, if there exists a monic $i \in \text{Hom}_Q((J, \mu_J), (J', \mu_{J'}))$ such that $j' \circ i = j$. A proper ideal $j : (J, \mu_J) \rightarrow (Q, m)$ in (Q, m, η) is called *maximal* if every proper ideal $j' : (J', \mu_{J'}) \rightarrow (Q, m)$ containing $j : (J, \mu_J) \rightarrow (Q, m)$ is isomorphic to $j : (J, \mu_J) \rightarrow (Q, m)$.

LEMMA 362. *Let \mathcal{D} be an essentially small abelian strict symmetric tensor category, (Q, m, η) a commutative monoid in \mathcal{D} . Then every proper ideal $j : (J, \mu_J) \rightarrow (Q, m)$ in (Q, m, η) is contained in a maximal ideal $\tilde{j} : (\tilde{J}, \tilde{\mu}) \rightarrow (Q, m)$.*

Proof. The ideals in (Q, m, η) do not necessarily form a set, but the isomorphism classes do, since \mathcal{D} is assumed essentially small. The relation $<$ on the ideals in (Q, m, η) gives rise to a partial ordering of the set of isomorphism classes of ideals. The maximal elements w.r.t. this partial order are precisely the isomorphism classes of maximal ideals. Now we can apply Zorn’s Lemma to complete the proof as in commutative algebra. ■

As in the category $R\text{-mod}$, we can quotient a commutative monoid by an ideal:

LEMMA 363. *Let \mathcal{D} be an abelian strict symmetric tensor category, (Q, m, η) a commutative monoid and $j : (J, \mu_J) \rightarrow (Q, m)$ an ideal. Let $p = \text{coker } j : (Q, m) \rightarrow (B, \mu_B)$. Then there exist unique morphisms $m_B : B \otimes B \rightarrow B$ and $\eta_B : \mathbf{1} \rightarrow B$ such that*

1. (B, m_B, η_B) is a commutative monoid,
2. $p \circ m = m_B \circ p \otimes p$,
3. $p \circ \eta = \eta_B$.

The monoid (B, m_B, η_B) is called the quotient of (Q, m, η) by the ideal $j : (J, \mu_J) \rightarrow (Q, m)$. It is nontrivial (B is not a zero object) iff the ideal is proper.

Furthermore, the map $p_\Gamma : \Gamma : \Gamma_Q \rightarrow \Gamma_B$ given by $s \mapsto p \circ s$ is a homomorphism of commutative algebras, which is surjective if the unit $\mathbf{1} \in \mathcal{D}$ is a projective object.

Proof. The construction of m_B, η_B proceeds essentially as in commutative algebra, despite the fact that the absence of elements makes it somewhat more abstract. Since $p : (Q, m) \rightarrow (B, \mu_B)$ is the cokernel of j , B is non-zero iff j is not epi, to wit if the ideal is proper. The equations $p \circ m = m_B \circ p \otimes p$ and $p \circ \eta = \eta_B$ imply that p_Γ is a unital homomorphism. If $\mathbf{1}$ is projective then the very Definition 352 implies that for every $s : \mathbf{1} \rightarrow B$ there is $t : \mathbf{1} \rightarrow Q$ such that $s = p \circ t$, thus p_Γ is surjective. ■

LEMMA 364. Let \mathcal{D} be an essentially small abelian strict symmetric tensor category. Let (Q, m, η) be a commutative monoid in \mathcal{D} and $j : (J, \mu) \rightarrow (Q, m)$ an ideal. Let (B, m_B, η_B) be the quotient monoid. Then there is a bijective correspondence between equivalence classes of ideals in (B, m_B, η_B) and equivalence classes of ideals $j' : (J', \mu') \rightarrow (Q, m)$ in (Q, μ, η) that contain $j : (J, \mu) \rightarrow (Q, m)$.

In particular, if j is a maximal ideal then all ideals in (B, m_B, η_B) are either zero or isomorphic to (B, m_B) .

Proof. As in ordinary commutative algebra. ■

LEMMA 365. Let k be a field and (Q, m, η) a commutative monoid in the strict symmetric abelian k -linear category \mathcal{D} . If every non-zero ideal in (Q, m, η) is isomorphic to (Q, m) then the commutative unital k -algebra $\text{End}_Q((Q, m))$ is a field.

Proof. Let $s \in \text{End}_Q((Q, m))$ be non-zero. Then $\text{im } s \neq 0$ is a non-zero ideal in (Q, m) , thus must be isomorphic to (Q, m) . Therefore $\text{im } s$ and in turn s are epi. Since $s \neq 0$, the kernel $\ker s$ is not isomorphic to (Q, m) and therefore it must be zero, thus s is monic. By Proposition 350, s is an isomorphism. Thus the commutative k -algebra $\text{End}_Q((Q, m))$ is a field extending k . ■

The following lemma is borrowed from [Bichon, 1998]:

LEMMA 366. Let \mathcal{D} be an abelian strict symmetric tensor category and (Q, m, η) a commutative monoid in it. Then every epimorphism in $\text{End}_Q((Q, m))$ is an isomorphism.

Proof. Let $g \in \text{End}_Q((Q, m))$ be an epimorphism and let $j : (J, \mu_J) \rightarrow (Q, m)$ be an ideal in (Q, m, η) . Now, $Q - \text{Mod}$ is a tensor category whose unit is (Q, m) , thus there is an isomorphism $s \in \text{Hom}_Q((J, \mu_J), (Q \otimes_Q J, \mu_{Q \otimes_Q J}))$. Let $h \in \text{End}_Q((J, \mu_J))$ be the composition

$$(J, \mu_J) \xrightarrow{s} (Q \otimes_Q J, \mu_{Q \otimes_Q J}) \xrightarrow{g \otimes \text{id}_J} (Q \otimes_Q J, \mu_{Q \otimes_Q J}) \xrightarrow{s^{-1}} (J, \mu_J).$$

Since the tensor product \otimes_Q of $Q - \text{Mod}$ is right-exact, $g \otimes \text{id}_J$ is epi. Now, $j \circ h = g \circ j$, and if we put $(j : (J, \mu_J) \rightarrow (Q, m)) = \ker g$ we have $j \circ h = 0$ and thus $j = 0$ since h is epi. Thus g is monic and therefore an isomorphism. ■

A.7 Inductive limits and the Ind-category

We need the categorical version of the concept of an inductive limit. For our purposes, inductive limits over \mathbb{N} will do, but in order to appeal to existing theories we need some definitions.

DEFINITION 367. If \mathcal{I}, \mathcal{C} are categories and $F : \mathcal{I} \rightarrow \mathcal{C}$ a functor, then a *colimit* (or *inductive limit*) of F consists of an object $Z \in \mathcal{C}$ and, for every $X \in \mathcal{I}$, of a morphism $i_X : F(X) \rightarrow Z$ in \mathcal{C} such that

1. $i_Y \circ F(s) = i_X$ for every morphism $s : X \rightarrow Y$ in \mathcal{I} .
2. Given $Z' \in \mathcal{C}$ and a family of morphisms $j_X : F(X) \rightarrow Z'$ in \mathcal{C} such that $j_Y \circ F(s) = j_X$ for every morphism $s : X \rightarrow Y$ in \mathcal{I} , there is a unique morphism $\iota : Z \rightarrow Z'$ such that $j_X = \iota \circ i_X$ for all $X \in \mathcal{I}$.

The second property required above is the universal property. It implies that any two colimits of F are isomorphic. Thus the colimit is essentially unique, provided it exists.

DEFINITION 368. A category \mathcal{I} is *filtered* if it is non-empty and

1. For any two objects $X, Y \in \mathcal{I}$ there is an $Z \in \mathcal{I}$ and morphisms $i : X \rightarrow Z, j : Y \rightarrow Z$.
2. For any two morphisms $u, v : X \rightarrow Y$ in \mathcal{I} there is a morphism $w : Y \rightarrow Z$ such that $w \circ u = w \circ v$.

Note that any directed partially ordered set (I, \leq) is a filtered category if we take the objects to be the elements of I , and the arrows are ordered pairs $\{(i, j) : i \leq j\}$.

DEFINITION 369. Let \mathcal{C} be a category. Then the category $\text{Ind } \mathcal{C}$ is defined as the functor category whose objects are all functors $F : \mathcal{I} \rightarrow \mathcal{C}$, where \mathcal{I} is a small filtered category. For $F : \mathcal{I} \rightarrow \mathcal{C}, F' : \mathcal{I}' \rightarrow \mathcal{C}$, the hom-set is defined by

$$\text{Hom}_{\text{Ind } \mathcal{C}}(F, F') = \varinjlim_x \varinjlim_y \text{Hom}_{\mathcal{C}}(F(X), F'(Y)).$$

(An element of the r.h.s. consists of a family $(f_{X,Y} : F(X) \rightarrow F'(Y))_{X \in \mathcal{I}, Y \in \mathcal{I}'}$ satisfying $F'(s) \circ f_{X,Y} = f_{X,Y'}$ for every $s : Y \rightarrow Y'$ in \mathcal{I}' and $f_{X',Y} \circ F(t) = f_{X,Y}$ for every $t : X \rightarrow X'$ in \mathcal{I} .) We leave it as an exercise to work out the composition of morphisms.

Some properties of $\text{Ind } \mathcal{C}$ are almost obvious. It contains \mathcal{C} as a subcategory: To every $X \in \mathcal{C}$ we assign the functor $F : \mathcal{I} \rightarrow \mathcal{C}$, where \mathcal{I} has only one object $*$ and $F(*) = X$. This embedding clearly is full and faithful. If \mathcal{C} is an Ab-category / additive / \mathbb{C} -linear then so is $\text{Ind } \mathcal{C}$. If \mathcal{C} is a strict (symmetric) tensor category then so is $\text{Ind } \mathcal{C}$: The tensor product of $F : \mathcal{I} \rightarrow \mathcal{C}$ and $F' : \mathcal{I}' \rightarrow \mathcal{C}$ is defined by $\mathcal{I}'' = \mathcal{I} \times \mathcal{I}'$ (which is a filtered category) and $F \otimes F' : \mathcal{I}'' \ni X \times Y \mapsto F(X) \otimes F'(Y)$. For the remaining results that we need, we just cite [SGA, 1972], to which we also refer for the proof:

THEOREM 370. *Ind \mathcal{C} has colimits for all small filtered index categories \mathcal{I} . If \mathcal{C} is an abelian category \mathcal{C} then $\text{Ind } \mathcal{C}$ is abelian.*

Thus every abelian (symmetric monoidal) category is a full subcategory of an abelian (symmetric monoidal) category that is complete under filtered colimits. For us this means that in $\text{Ind } \mathcal{C}$ we can make sense of infinite direct sums indexed by \mathbb{N} , defining $\bigoplus_{i \in \mathbb{N}} X_i$ as the colimit of the functor $F : \mathcal{I} \rightarrow \mathcal{C}$, where \mathcal{I} is the

poset \mathbb{N} interpreted as a filtered category, and $F(n) = \bigoplus_{i=1}^n X_i$ together with the obvious morphisms $F(n) \rightarrow F(m)$ when $n \leq m$.

LEMMA 371. *If \mathcal{C} is a TC^* then every object $X \in \mathcal{C}$ is projective as an object of $\text{Ind } \mathcal{C}$.*

Proof. First assume that X is irreducible and consider $s : X \rightarrow B$. Given an epi $p : A \rightarrow B$ in $\text{Ind } \mathcal{C}$, we have $A = \lim_{\leftarrow} A_i$ with $A_i \in \mathcal{C}$ and similarly for B . Furthermore, $\text{Hom}(A, B) = \lim_{\leftarrow} \lim_{\leftarrow} \text{Hom}_{\mathcal{C}}(A_i, B_j)$ and $\text{Hom}(X, B) = \lim_{\leftarrow} \text{Hom}_{\mathcal{C}}(X, B_j)$. Since X is irreducible and \mathcal{C} is semisimple, X is a direct summand of B_j whenever $s_j : X \rightarrow B_j$ is non-zero. Since $p : A \rightarrow B$ is epi, the component $A_i \rightarrow B_j$ is epi for i sufficiently big. By semisimplicity of \mathcal{C} , s_j then lifts to a morphism $X \rightarrow A_i$. Putting everything together this gives a morphism $\widehat{s} : X \rightarrow A$ such that $p \circ \widehat{s} = s$.

Now let X be a finite direct sum of irreducible X_i with isometries $v_i : X_i \rightarrow X$ and $s : X \rightarrow B$. Defining $s_i = s \circ v_i : X_i \rightarrow B$, the first half of the proof provides $\widehat{s}_i : X_i \rightarrow A$ such that $p \circ \widehat{s}_i = s_i$. Now define $\widehat{s} = \sum_i \widehat{s}_i \circ v_i^* : X \rightarrow A$. We have

$$p \circ \widehat{s} = \sum_i p \circ \widehat{s}_i \circ v_i^* = \sum_i s_i \circ v_i^* = \sum_i s \circ v_i \circ v_i^* = s,$$

proving projectivity of X . ■

B ABSTRACT DUALITY THEORY FOR SYMMETRIC TENSOR *-CATEGORIES

In the first two subsections we give self-contained statements of the results needed for the reconstructions in AQFT. Some of the proofs are deferred to the rest of this appendix, which hurried (or less ambitious) or readers may safely skip.

B.1 Fiber functors and the concrete Tannaka theorem. Part I

Let $\text{Vect}_{\mathbb{C}}$ denote the \mathbb{C} -linear symmetric tensor category of finite dimensional \mathbb{C} -vector spaces and \mathcal{H} denote the STC^* of finite dimensional Hilbert spaces. We pretend that both tensor categories are strict, which amounts to suppressing the associativity and unit isomorphisms α, λ, ρ from the notation. Both categories have a canonical symmetry Σ , the flip isomorphism $\Sigma_{V, V'} : V \otimes V' \rightarrow V' \otimes V$.

DEFINITION 372. Let \mathcal{C} be an STC^* . A fiber functor for \mathcal{C} is a faithful \mathbb{C} -linear tensor functor $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$. A $*$ -preserving fiber functor for \mathcal{C} is a faithful functor $E : \mathcal{C} \rightarrow \mathcal{H}$ of tensor $*$ -categories. E is symmetric if $E(c_{X, Y}) = \Sigma_{E(X), E(Y)}$, i.e. the symmetry of \mathcal{C} is mapped to the canonical symmetry of $\text{Vect}_{\mathbb{C}}$ or \mathcal{H} , respectively.

An STC^* equipped with a symmetric $*$ -preserving fiber functor is called concrete, since it is equivalent to a (non-full!) tensor subcategory of the category \mathcal{H} of Hilbert spaces. Our main concern in this appendix are (1) Consequences of the

existence of a fiber functor, (2) Uniqueness of fiber functors, and (3) Existence of fiber functors. As to (2) we will prove:

THEOREM 373. *Let \mathcal{C} be an STC^* and let $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{H}$ be $*$ -preserving symmetric fiber functors. Then $E_1 \cong E_2$, i.e. there exists a unitary monoidal natural isomorphism $\alpha : E_1 \rightarrow E_2$.*

We now *assume* a symmetric $*$ -preserving fiber functor for the STC^* \mathcal{C} to be given. Let $G_E \subset \text{Nat}_{\otimes} E$ denote the set of unitary monoidal natural transformations of E (to itself). This clearly is a group with the identical natural transformation as unit. G_E can be identified with a subset of $\prod_{X \in \mathcal{C}} \mathcal{U}(E(X))$, where $\mathcal{U}(E(X))$ is the compact group of unitaries on the finite dimensional Hilbert space $E(X)$. The product of these groups is compact by Tychonov's theorem, cf. e.g. [Pedersen, 1989, Theorem 1.6.10], and since G_E is a closed subset, it is itself compact. The product and inverse maps are continuous, thus G_E is a compact topological group. By its very definition, the group G_E acts on the Hilbert spaces $E(X), X \in \mathcal{C}$ by unitary representations π_X , namely $\pi_X(g) = g_X$ where g_X is the component at X of the natural transformation $g \in G_E$.

PROPOSITION 374. *There is a faithful symmetric tensor $*$ -functor $F : \mathcal{C} \rightarrow \text{Rep}_f G_E$ such that $K \circ F = E$, where $K : \text{Rep}_f G_E \rightarrow \mathcal{H}$ is the forgetful functor $(H, \pi) \mapsto H$.*

Proof. We define $F(X) = (E(X), \pi_X) \in \text{Rep}_f G_E$ for all $X \in \mathcal{C}$ and $F(s) = E(s)$ for all $s \in \text{Hom}(X, Y)$. For $s : X \rightarrow Y$ we have

$$F(s)\pi_X(g) = F(s)g_X = g_Y F(s) = \pi_Y(g)F(s),$$

since $g : E \rightarrow E$ is a natural transformation. Thus F is a functor, which is obviously $*$ -preserving and faithful. In view of $g_{\mathbf{1}} = \text{id}_{E(\mathbf{1})}$ for every $g \in G_E$, we have $F(\mathbf{1}_{\mathcal{C}}) = (\mathbb{C}, \pi_0) = \mathbf{1}_{\text{Rep}_f G_E}$, where π_0 is the trivial representation. In order to see that F is a functor of tensor $*$ -categories we must produce unitaries $d_{X,Y}^F : F(X) \otimes F(Y) \rightarrow F(X \otimes Y), X, Y \in \mathcal{C}$ and $e : \mathbf{1}_{\text{Rep}_f G_E} \rightarrow F(\mathbf{1}_{\mathcal{C}})$ satisfying (56) and (57), respectively. We claim that the choice $e^F = e^E, d_{X,Y}^F = d_{X,Y}^E$ does the job, where the e^E and $d_{X,Y}^E$ are the unitaries coming with the tensor functor $E : \mathcal{C} \rightarrow \mathcal{H}$. It is obvious that e^E and $d_{X,Y}^E$ satisfy (56) and (57), but we must show that they are morphisms in $\text{Rep}_f G_E$. For $d_{X,Y}^E$ this follows from the computation

$$d_{X,Y}^F \circ (\pi_X(g) \otimes \pi_Y(g)) = d_{X,Y}^E \circ g_X \otimes g_Y = g_{X \otimes Y} \circ d_{X,Y}^E = \pi_{X \otimes Y}(g) \circ d_{X,Y}^F,$$

where we have used that g is a monoidal natural transformation. Now, by the definition of a natural monoidal transformation we have $g_{\mathbf{1}} = \text{id}_{E(\mathbf{1})}$ for all $g \in G_E$, i.e. $F(\mathbf{1}) = (E(\mathbf{1}), \pi_{\mathbf{1}})$ is the trivial representation. If the strict unit $\mathbf{1}_{\mathcal{H}} = \mathbb{C}$ is in the image of E then, by naturality, it also carries the trivial representation, thus e^F in fact is a morphism of representations. (In case $\mathbf{1}_{\mathcal{H}} \notin E(\mathcal{C})$, we equip $\mathbf{1}_{\mathcal{H}}$ with the trivial representation by hand.) Since the symmetry of $\text{Rep}_f G_E$ is by definition given by $c((H, \pi), (H', \pi')) = c(H, H')$, where the right hand side refers

to the category \mathcal{H} , and since E respects the symmetries, so does F . $K \circ F = E$ is obvious. ■

The proof of the following proposition is postponed, since it requires further preparations.

PROPOSITION 375. *Let \mathcal{C} be an STC^* and $E : \mathcal{C} \rightarrow \mathcal{H}$ a symmetric $*$ -preserving fiber functor. Let G_E and $F : \mathcal{C} \rightarrow \text{Rep}_f G_E$ as defined above. Then the following hold:*

- (i) *If $X \in \mathcal{C}$ is irreducible then $\text{span}_{\mathbb{C}}\{\pi_X(g), g \in G_E\}$ is dense in $\text{End } E(X)$.*
- (ii) *If $X, Y \in \mathcal{C}$ are irreducible and $X \not\cong Y$ then $\text{span}_{\mathbb{C}}\{\pi_X(g) \oplus \pi_Y(g), g \in G_E\}$ is dense in $\text{End } E(X) \oplus \text{End } E(Y)$.*

THEOREM 376. *Let \mathcal{C} be an STC^* and $E : \mathcal{C} \rightarrow \mathcal{H}$ a symmetric $*$ -preserving fiber functor. Let G_E and $F : \mathcal{C} \rightarrow \text{Rep}_f G_E$ as defined above. Then F is an equivalence of symmetric tensor $*$ -categories.*

Proof. We already know that F is a faithful symmetric tensor functor. In view of Proposition 318 it remains to show that F is full and essentially surjective.

Since the categories \mathcal{C} and $\text{Rep}_f G_E$ are semisimple, in order to prove that F is full it is sufficient to show that (a) $F(X) \in \text{Rep}_f G_E$ is irreducible if $X \in \mathcal{C}$ is irreducible and (b) whenever $X, Y \in \mathcal{C}$ are irreducible and inequivalent then $\text{Hom}(F(X), F(Y)) = \{0\}$. Now, (i) of Proposition 375 clearly implies that $\text{End}(F(X)) = \mathbb{C}\text{id}$, which is the desired irreducibility of $F(X)$. Assume now that $X, Y \in \mathcal{C}$ are irreducible and non-isomorphic and let $s \in \text{Hom}(F(X), F(Y))$, to wit $s \in \text{Hom}(E(X), E(Y))$ and $s\pi_X(g) = \pi_Y(g)s$ for all $g \in G_E$. Then (ii) of Proposition 375 implies $su = vs$ for any $u \in \text{End } E(X)$ and $v \in \text{End } E(Y)$. With $u = 0$ and $v = 1$ this implies $s = 0$, thus the irreps $F(X) = (E(X), \pi_X)$ and $F(Y) = (E(Y), \pi_Y)$ are non-isomorphic. This proves that F is full.

Therefore, F is an equivalence of \mathcal{C} with a full tensor subcategory of $\text{Rep}_f G_E$. If $g \in G_E$ is nontrivial, it is immediate by the definition of G_E that there is an $X \in \mathcal{C}$ such that $g_X \neq \text{id}_{E(X)}$ – but this means $\pi_X(g) \neq 1$. In other words, the representations $\{F(X), X \in \mathcal{C}\}$ separate the points of G_E . But it is a well known consequence of the Peter-Weyl theorem that a full monoidal subcategory of $\text{Rep}_f G_E$ separates the points of G_E iff it is in fact equivalent to $\text{Rep}_f G_E$. Thus the functor F is essentially surjective, and we are done. ■

Since they are so important, we restate Theorems 373 and 376 in a self contained way:

THEOREM 377. *Let \mathcal{C} be an STC^* and $E : \mathcal{C} \rightarrow \mathcal{H}$ a $*$ -preserving symmetric fiber functor. Let G_E be the group of unitary monoidal natural transformations of E with the topology inherited from $\prod_{X \in \mathcal{C}} \mathcal{U}(E(X))$. Then G_E is compact and the functor $F : \mathcal{C} \rightarrow \text{Rep}_f G_E, X \mapsto (E(X), \pi_X)$, where $\pi_X(g) = g_X$, is an equivalence*

of STC^* s. If $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{H}$ are $*$ -preserving symmetric fiber functors then $E_1 \cong E_2$ and therefore $G_{E_1} \cong G_{E_2}$.

REMARK 378. The preceding theorem is essentially a reformulation in modern language of the classical result of Tannaka [Tannaka, 1939]. It can be generalized, albeit without the uniqueness part, to a setting where \mathcal{C} is only braided or even has no braiding. This leads to a (concrete) Tannaka theory for quantum groups, for which the interested reader is referred to the reviews [Joyal and Street, 1991] and [Müger *et al.*, 2004].

Before we turn to proving Theorem 373 (Subsection B.4) and Proposition 375 (Subsection B.5) we identify a necessary condition for the existence of fiber functors, which will lead us to a generalization of Theorem 377.

B.2 Compact supergroups and the abstract Tannaka theorem

According to Theorem 377, an STC^* admitting a symmetric $*$ -preserving fiber functor is equivalent, as a symmetric tensor $*$ -category, to the category of finite dimensional unitary representations of a compact group G that is uniquely determined up to isomorphism. Concerning the existence of fiber functors it will turn out that the twist Θ (Definition 342) provides an obstruction, fortunately the only one.

DEFINITION 379. An STC^* is called even if $\Theta(X) = \text{id}_X$ for all $X \in \mathcal{C}$.

EXAMPLE 380. A simple computation using the explicit formulae for $r, \bar{r}, c_{X,Y}$ given in Example 333 shows that the STC^* \mathcal{H} of finite dimensional Hilbert spaces is even. The same holds for the category $\text{Rep}_f G$ of finite dimensional unitary representations of a compact group G .

This suggests that an STC^* must be even in order to admit a fiber functor. In fact:

PROPOSITION 381. *If an STC^* \mathcal{C} admits a $*$ -preserving symmetric fiber functor E then it is even.*

Proof. By Proposition 344, we have $E(\Theta(X)) = \Theta(E(X))$. Since \mathcal{H} is even, this equals $\text{id}_{E(X)} = E(\text{id}_X)$. Since E is faithful, this implies $\Theta(X) = \text{id}_X$. ■

Fortunately, this is the only obstruction since, beginning in the next subsection, we will prove:

THEOREM 382. *Every even STC^* admits a $*$ -preserving symmetric fiber functor $E : \mathcal{C} \rightarrow \mathcal{H}$.*

Combining this with Theorem 377 we obtain:

THEOREM 383. *Let \mathcal{C} be an even STC^* . Then there is a compact group G , unique up to isomorphism, such that there exists an equivalence $F : \mathcal{C} \rightarrow \text{Rep}_f G$ of STC^* s.*

Theorem 383 is not yet sufficiently general for the application to quantum field theory, which is the subject of this chapter. Making the connection with DHR theory, we see that the twist of an irreducible DHR sector is ± 1 , depending on whether the sector is bosonic or fermionic. Since in general we cannot a priori rule out fermionic sectors, we cannot restrict ourselves to even STC^* s. What we therefore really need is a characterization of all STC^* s. This requires a generalization of the notion of compact groups:

DEFINITION 384. A (compact) supergroup is a pair (G, k) where G is a (compact Hausdorff) group and k is an element of order two in the center of G . An isomorphism $\alpha : (G, k) \xrightarrow{\cong} (G', k')$ of (compact) supergroups is an isomorphism $\alpha : G \rightarrow G'$ of (topological) groups such that $\alpha(k) = k'$.

DEFINITION 385. A (finite dimensional, unitary, continuous) representation of a compact supergroup (G, k) is just a (finite dimensional, unitary, continuous) representation (H, π) of G . Intertwiners and the tensor product of representations are defined as for groups, thus $\text{Rep}_{(f)}(G, k) \cong \text{Rep}_{(f)}G$ as C^* -tensor categories. (Since k is in the center of G , morphisms in $\text{Rep}_{(f)}(G, k)$ automatically preserve the \mathbb{Z}_2 -grading induced by $\pi(k)$. $\text{Rep}_{(f)}(G, k)$ is equipped with a symmetry Σ_k as follows: For every $(H, \pi) \in \text{Rep}(G, k)$ let $P_{\pm} = (\text{id} + \pi(k))/2$ be the projector on the even and odd subspaces of a representation space H , respectively. Then

$$\Sigma_k((H, \pi), (H', \pi')) = \Sigma(H, H')(1 - 2P_{-}^{\pi} \otimes P_{-}^{\pi'}),$$

where $\Sigma(H, H') : H \otimes H' \rightarrow H' \otimes H$ is the usual flip isomorphism $x \otimes y \mapsto y \otimes x$. Thus for homogeneous $x \in H, y \in H'$ we have $\Sigma_k((H, \pi), (H', \pi')) : x \otimes y \mapsto \pm y \otimes x$, where the minus sign occurs iff $x \in H_{-}$ and $y \in H'_{-}$. In the case $(G, k) = (\{e, k\}, k)$, we call $\text{Rep}_f(G, k)$ the category \mathcal{SH} of super Hilbert spaces.

REMARK 386. Note that the action of k induces a \mathbb{Z}_2 -grading on H that is stable under the G -action. Since the symmetry Σ_k defined above is precisely the one on the category \mathcal{SH} of finite dimensional super Hilbert spaces, we see that there is a forgetful symmetric tensor functor $\text{Rep}_f(G, k) \rightarrow \mathcal{SH}$.

LEMMA 387. Σ_k as defined above is a symmetry on the category $\text{Rep}(G, k)$. Thus $\text{Rep}_f(G, k)$ is a STC^* . For every object $X = (H, \pi) \in \text{Rep}_f(G, k)$, the twist $\Theta(X)$ is given by $\pi(k)$.

Proof. Most of the claimed properties follow immediately from those of Rep_fG . It is clear that $\Sigma_k((H, \pi), (H', \pi')) \circ \Sigma_k((H', \pi'), (H, \pi))$ is the identity of $H' \otimes H$. We only need to prove naturality and compatibility with the tensor product. This is an easy exercise. The same holds for the identity $\Theta((H, \pi)) = \pi(k)$. ■

We need a corollary of (the proof of) Theorem 383:

COROLLARY 388. For any compact group G , the unitary monoidal natural transformations of the identity functor on Rep_fG form an abelian group that is isomorphic to the center $Z(G)$.

Proof. If $k \in Z(G)$ and $(H, \pi) \in \text{Rep}_f G$ is irreducible then $\pi(k) = \omega_{(H, \pi)} \text{id}_H$, where $\omega_{(H, \pi)}$ is a scalar. Defining $\Theta((H, \pi)) = \omega_{(H, \pi)} \text{id}_{(H, \pi)}$ and extending to reducible objects defines a unitary monoidal natural isomorphism of $\text{Rep}_f G$. Conversely, let $\{\Theta((H, \pi))\}$ be a unitary monoidal isomorphism of the identity functor of $\text{Rep}_f G$ and $K : \text{Rep}_f G \rightarrow \mathcal{H}$ the forgetful functor. Then the family $(\alpha_{(H, \pi)} = K(\Theta((H, \pi))))$ is a unitary monoidal natural isomorphism of K . By Theorem 377, there is a $g \in G$ such that $\alpha_{(H, \pi)} = \pi(g)$ for all $(H, \pi) \in \text{Rep}_f G$. Since $\pi(g)$ is a multiple of the identity for every irreducible (H, π) , g is in $Z(G)$ by Schur’s lemma. Clearly the above correspondence is an isomorphism of abelian groups. ■

Modulo Theorem 382 we can now prove the Main Result of this appendix:

THEOREM 389. *Let \mathcal{C} be an STC^* . Then there exist a compact supergroup (G, k) , unique up to isomorphism, and an equivalence $F : \mathcal{C} \rightarrow \text{Rep}_f(G, k)$ of symmetric tensor $*$ -categories. In particular, if $K : \text{Rep}_f(G, k) \rightarrow \mathcal{SH}$ is the forgetful functor, the composite $E = K \circ F : \mathcal{C} \rightarrow \mathcal{SH}$ is a ‘super fiber functor’, i.e. a faithful symmetric $*$ -preserving tensor functor into the STC^* of super Hilbert spaces.*

Proof. We define a new STC^* $\tilde{\mathcal{C}}$ (the ‘bosonization’ of \mathcal{C}) as follows. As a tensor $*$ -category, $\tilde{\mathcal{C}}$ coincides with \mathcal{C} . The symmetry \tilde{c} is defined by

$$\tilde{c}_{X, Y} = (-1)^{(1-\Theta(X))(1-\Theta(Y))/4} c_{X, Y}$$

for irreducible $X, Y \in \text{Obj } \mathcal{C} = \text{Obj } \tilde{\mathcal{C}}$, and extended to all objects by naturality. It is easy to verify that $(\tilde{\mathcal{C}}, \tilde{c})$ is again a symmetric tensor category, in fact an even one. Thus by Theorem 383 there is a compact group G such that $\tilde{\mathcal{C}} \simeq \text{Rep}_f G$ as STC^* s. Applying Corollary 388 to the category $\tilde{\mathcal{C}} \simeq \text{Rep}_f G$ and the family $(\Theta(X))_{X \in \mathcal{C}}$, as defined in the original category \mathcal{C} proves the existence of an element $k \in Z(G), k^2 = e$, such that $\Theta((H, \pi)) = \pi(k)$ for all $(H, \pi) \in \tilde{\mathcal{C}} \simeq \text{Rep}_f G$. Clearly (G, k) is a supergroup. We claim that $\mathcal{C} \simeq \text{Rep}_f(G, k)$ as STC^* s. Ignoring the symmetries this is clearly true since $\text{Rep}_f(G, k) \simeq \text{Rep}_f G$ as tensor $*$ -categories. That \mathcal{C} and $\text{Rep}_f(G, k)$ are equivalent as STC^* s, i.e. taking the symmetries into account, follows from the fact that \mathcal{C} is related to $\tilde{\mathcal{C}}$ precisely as $\text{Rep}_f(G, k)$ is to $\text{Rep}_f G$, namely by a twist of the symmetry effected by the family $(\Theta((H, \pi)) = \pi(k))$. To conclude, we observe that the uniqueness result for (G, k) follows from the uniqueness of G in Theorem 383 and that of k in Corollary 388. ■

REMARK 390. Theorem 389 was proven by Doplicher and Roberts in [Doplicher and Roberts, 1989, Section 7] exactly as stated above, the only superficial difference being that the terminology of supergroups wasn’t used. (Note that our supergroups are not what is usually designated by this name.) As above, the proof was by

reduction to even categories and compact groups. Independently and essentially at the same time, a result analogous to Theorem 382 for (pro)algebraic groups was proven by Deligne in [Deligne, 1990], implying an algebraic analogue of Theorem 383 by [Saavedra Rivano, 1972; Deligne and Milne, 1982]. Recently, Deligne also discussed the super case, cf. [Deligne, 2002].

This concludes the discussion of the main results of this appendix. We now turn to proving Theorem 373, Proposition 375 and Theorem 382.

B.3 Certain algebras arising from fiber functors

Let \mathcal{C} be a TC^* and $E_1, E_2 : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ fiber functors. Recall that they come with natural isomorphisms $d_{X,Y}^i : E_i(X) \otimes E_i(Y) \rightarrow E_i(X \otimes Y)$ and $e^i : \mathbf{1}_{\text{Vect}} = \mathbb{C} \rightarrow E_i(\mathbf{1}_{\mathcal{C}})$. Consider the \mathbb{C} -vector space

$$A_0(E_1, E_2) = \bigoplus_{X \in \mathcal{C}} \text{Hom}(E_2(X), E_1(X)).$$

For $X \in \mathcal{C}$ and $s \in \text{Hom}(E_2(X), E_1(X))$ we write $[X, s]$ for the element of $A_0(E_1, E_2)$ which takes the value s at X and is zero elsewhere. Clearly, A_0 consists precisely of the finite linear combinations of such elements. We turn $A_0(E_1, E_2)$ into a \mathbb{C} -algebra by defining $[X, s] \cdot [Y, t] = [X \otimes Y, u]$, where u is the composite

$$E_2(X \otimes Y) \xrightarrow{(d_{X,Y}^2)^{-1}} E_2(X) \otimes E_2(Y) \xrightarrow{s \otimes t} E_1(X) \otimes E_1(Y) \xrightarrow{d_{X,Y}^1} E_1(X \otimes Y)$$

Since \mathcal{C} is strict, we have $(X \otimes Y) \otimes Z = X \otimes (Y \otimes Z)$ and $\mathbf{1} \otimes X = X = X \otimes \mathbf{1}$. Together with the 2-cocycle type equation (56) satisfied by the isomorphisms $d_{X,Y}^i$ this implies that $A_0(E_1, E_2)$ is associative. The compatibility (57) of $d_{X,Y}^i$ with e^i for $i = 1, 2$ implies that $[\mathbf{1}, e^1 \circ (e^2)^{-1}]$ is a unit of the algebra $A_0(E_1, E_2)$.

LEMMA 391. *Let \mathcal{C} be a TC^* and $E_1, E_2 : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ fiber functors. The subspace*

$$I(E_1, E_2) = \text{span}_{\mathbb{C}}\{[X, a \circ E_2(s)] - [Y, E_1(s) \circ a] \mid s : X \rightarrow Y, a : E_2(Y) \rightarrow E_1(X)\}$$

is a two-sided ideal.

Proof. To show that $I(E_1, E_2) \subset A_0(E_1, E_2)$ is an ideal, let $s : X \rightarrow Y$, $a \in \text{Hom}(E_2(Y), E_1(X))$, thus $[X, a \circ E_2(s)] - [Y, E_1(s) \circ a] \in I(E_1, E_2)$, and let $[Z, t] \in A_0(E_1, E_2)$. Then

$$\begin{aligned} & ([X, a \circ E_2(s)] - [Y, E_1(s) \circ a]) \cdot [Z, t] \\ &= [X \otimes Z, d_{X,Z}^1 \circ (a \circ E_2(s)) \otimes t \circ (d_{X,Z}^2)^{-1}] \\ &\quad - [Y \otimes Z, d_{Y,Z}^1 \circ (E_1(s) \circ a) \otimes t \circ (d_{Y,Z}^2)^{-1}] \\ &= [X \otimes Z, d_{X,Z}^1 \circ a \otimes t \circ (d_{Y,Z}^2)^{-1} \circ E_2(s \otimes \text{id}_Z)] \\ &\quad - [Y \otimes Z, E_1(s \otimes \text{id}_Z) \circ d_{X,Z}^1 \circ a \otimes t \circ (d_{Y,Z}^2)^{-1}] \\ &= [X', a' \circ E_2(s')] - [Y', E_1(s') \circ a'] \in I(E_1, E_2), \end{aligned}$$

where in the second equality we used naturality of d^i , and in the last line we wrote $X' = X \otimes Z, Y' = Y \otimes Z, s' = s \otimes \text{id}_Z : X' \rightarrow Y'$ and $a' = d_{X,Z}^1 \circ a \otimes t \circ (d_{Y,Z}^2)^{-1} \in \text{Hom}(E_2(Y'), E_1(X'))$ in order to make clear that the result is in $I(E_1, E_2)$. This proves that the latter is a left ideal in $A_0(E_1, E_2)$. Similarly, one shows that it is a right ideal. ■

We denote by $A(E_1, E_2)$ the quotient algebra $A_0(E_1, E_2)/I(E_1, E_2)$. It can also be understood as the algebra generated by symbols $[X, s]$, where $X \in \mathcal{C}, s \in \text{Hom}(E_2(X), E_1(X))$, subject to the relations $[X, s] + [X, t] = [X, s + t]$ and $[X, a \circ E_2(s)] = [Y, E_1(s) \circ a]$ whenever $s : X \rightarrow Y, a \in \text{Hom}(E_2(Y), E_1(X))$. Therefore it should not cause confusion that we denote the image of $[X, s] \in A_0(E_1, E_2)$ in $A(E_1, E_2)$ again by $[X, s]$.

PROPOSITION 392. *Let \mathcal{C} be an STC* and $E_1, E_2 : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ fiber functors. If E_1, E_2 are symmetric then $A(E_1, E_2)$ is commutative.*

Proof. Assume that \mathcal{C} is symmetric and that the fiber functors satisfy $E_i(c_{X,Y}) = \Sigma_{E_i(X), E_i(Y)}$. Let $[A, u], [B, v] \in A_0(E_1, E_2)$, thus $A, B \in \mathcal{C}$ and $u : E_2(A) \rightarrow E_1(A), v : E_2(B) \rightarrow E_1(B)$. Then

$$[A, u] \cdot [B, v] = [A \otimes B, d_{A,B}^1 \circ u \otimes v \circ (d_{A,B}^2)^{-1}]$$

and

$$\begin{aligned} [B, v] \cdot [A, u] &= [B \otimes A, d_{B,A}^1 \circ v \otimes u \circ (d_{B,A}^2)^{-1}] \\ &= [B \otimes A, d_{B,A}^1 \circ \Sigma_{E_1(A), E_2(B)} \circ u \otimes v \circ \Sigma_{E_2(B), E_1(A)} \circ (d_{B,A}^2)^{-1}] \\ &= [B \otimes A, d_{B,A}^1 \circ E_1(c_{B,A}) \circ u \otimes v \circ E_2(c_{B,A}) \circ (d_{B,A}^2)^{-1}] \\ &= [B \otimes A, E_1(c_{A,B}) \circ d_{A,B}^1 \circ u \otimes v \circ (d_{A,B}^2)^{-1} \circ E_2(c_{B,A})]. \end{aligned}$$

With $X = A \otimes B, Y = B \otimes A, s = c_{A,B}$ and $a = d_{A,B}^1 \circ u \otimes v \circ (d_{A,B}^2)^{-1} \circ E_2(c_{B,A})$ we obtain

$$[A, u] \cdot [B, v] = [X, a \circ E_2(s)],$$

$$[B, v] \cdot [A, u] = [Y, E_1(s) \circ a].$$

Thus

$$[A, u] \cdot [B, v] - [B, v] \cdot [A, u] = [X, a \circ E_2(s)] - [Y, E_1(s) \circ a] \in I(E_1, E_2),$$

implying $[A_0(E_1, E_2), A_0(E_1, E_2)] \subset I(E_1, E_2)$. This implies commutativity of $A(E_1, E_2) = A_0(E_1, E_2)/I(E_1, E_2)$. ■

PROPOSITION 393. *Let \mathcal{C} be a TC* and let $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{H}$ be *-preserving fiber functors. Then $A(E_1, E_2)$ has a positive *-operation, i.e. an antilinear and antimultiplicative involution such that $a^*a = 0$ implies $a = 0$.*

Proof. We define a $*$ -operation \star on $A_0(E_1, E_2)$. Let $[X, s] \in A_0(E_1, E_2)$. Pick a standard conjugate $(\bar{X}_i, r_i, \bar{r}_i)$, define $t \in \text{Hom}_{\mathcal{H}}(E_2(\bar{X}), E_1(\bar{X}))$ by

$$t = \text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}^*) \circ \text{id}_{E_1(\bar{X})} \otimes s^* \otimes \text{id}_{E_2(\bar{X})} \circ E_1(r) \otimes \text{id}_{E_2(\bar{X})}$$

and put $[X, s]^* := [\bar{X}, t]$. (Of course, s^* is defined using the inner products on the Hilbert spaces $E_1(X), E_2(X)$.) If we pick another standard conjugate (\bar{X}', r', \bar{r}') of X , we know that there is a unitary $u : \bar{X} \rightarrow \bar{X}'$ such that $r' = u \otimes \text{id}_X \circ r$ and $\bar{r}' = \text{id}_X \otimes u \circ \bar{r}$. Using (\bar{X}', r', \bar{r}') we obtain $([X, s]^*)' := [\bar{X}', t']$ with t' defined by replacing r, \bar{r} by r', \bar{r}' . Now, $[\bar{X}, t] - [\bar{X}', t']$ equals

$$\begin{aligned} & [\bar{X}, \text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}^*) \circ \text{id}_{E_1(\bar{X})} \otimes s^* \otimes \text{id}_{E_2(\bar{X})} \circ E_1(r) \otimes \text{id}_{E_2(\bar{X})}] \\ & - [\bar{X}', \text{id}_{E_1(\bar{X}')} \otimes E_2(\bar{r}'^*) \circ \text{id}_{E_1(\bar{X}')} \otimes s^* \otimes \text{id}_{E_2(\bar{X}')} \circ E_1(r') \otimes \text{id}_{E_2(\bar{X}')}] \\ & = [\bar{X}, (\text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}'^*) \circ \text{id}_{E_1(\bar{X})} \otimes s^* \otimes \text{id}_{E_2(\bar{X}')} \circ E_1(r) \otimes \text{id}_{E_2(\bar{X}')}) \circ E_2(u)] \\ & - [\bar{X}', E_1(u) \circ (\text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}'^*) \circ \text{id}_{E_1(\bar{X})} \otimes s^* \otimes \text{id}_{E_2(\bar{X}')}) \circ E_1(r) \otimes \text{id}_{E_2(\bar{X}')}], \end{aligned}$$

which is in the ideal $I(E_1, E_2)$ defined in Proposition 398. Thus, while $[X, s]^*$ depends on the chosen conjugate (\bar{X}, r, \bar{r}) of X , its image $\gamma([X, s]^*) \in A(E_1, E_2)$ doesn't.

In order to be able to define a $*$ -operation on $A(E_1, E_2)$ by $x^* := \gamma \circ \star \circ \gamma^{-1}(x)$ we must show that the composite map $\gamma \circ \star : A_0(E_1, E_2) \rightarrow A(E_1, E_2)$ maps $I(E_1, E_2)$ to zero. To this purpose, let $X, Y \in \mathcal{C}, s : X \rightarrow Y, a \in \text{Hom}(E_2(Y), E_1(X))$ and choose conjugates $(\bar{X}, r_X, \bar{r}_X), (\bar{Y}, r_Y, \bar{r}_Y)$. Then

$$\begin{aligned} & [X, a \circ E_2(s)]^* - [Y, E_1(s) \circ a]^* \\ & = [\bar{X}, \text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}_X^*) \circ \text{id}_{E_1(\bar{X})} \otimes (a \circ E_2(s))^* \otimes \text{id}_{E_2(\bar{X})} \circ E_1(r_X) \otimes \text{id}_{E_2(\bar{X})}] \\ & - [\bar{Y}, \text{id}_{E_1(\bar{Y})} \otimes E_2(\bar{r}_Y^*) \circ \text{id}_{E_1(\bar{Y})} \otimes (E_1(s) \circ a)^* \otimes \text{id}_{E_2(\bar{Y})} \circ E_1(r_Y) \otimes \text{id}_{E_2(\bar{Y})}] \\ & = [\bar{X}, \tilde{a} \circ E_2(\tilde{s})] - [\bar{Y}, E_1(\tilde{s}) \circ \tilde{a}], \end{aligned}$$

with $\tilde{a} \in \text{Hom}_{\mathcal{H}}(E_2(\bar{Y}), E_1(\bar{X}))$ and $\tilde{s} \in \text{Hom}(\bar{X}, \bar{Y})$ defined by

$$\begin{aligned} \tilde{a} &= \text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}_X^*) \circ \text{id}_{E_1(\bar{X})} \otimes a^* \otimes \text{id}_{E_2(\bar{Y})} \circ E_1(r_X) \otimes \text{id}_{E_2(\bar{Y})}, \\ \tilde{s} &= \text{id}_{\bar{Y}} \otimes \bar{r}_X^* \circ \text{id}_{\bar{Y}} \otimes s^* \otimes \text{id}_{\bar{X}} \circ r_Y \otimes \text{id}_{\bar{X}}. \end{aligned}$$

This clearly is in $I(E_1, E_2)$, thus $x^* := \gamma \circ \star \circ \gamma^{-1}(x)$ defines a $*$ -operation on $A(E_1, E_2)$.

Now it is obvious that the resulting map $*$ on $A(E_1, E_2)$ is additive and antilinear. It also is involutive and antimultiplicative as one verifies by an appropriate use of the conjugate equations. We omit the tedious but straightforward computations. It remains to show positivity of the $*$ -operation. Consider $[X, s] \in A_0(E_1, E_2)$, pick a conjugate (\bar{X}, r, \bar{r}) and compute $[X, s]^* \cdot [X, s] = [\bar{X} \otimes X, t]$, where t equals

$$d_{\bar{X}, X}^1 \circ \left(\text{id}_{E_1(\bar{X})} \otimes E_2(\bar{r}^*) \circ \text{id}_{E_1(\bar{X})} \otimes s^* \otimes \text{id}_{E_2(\bar{X})} \circ E_1(r) \otimes \text{id}_{E_2(\bar{X})} \right) \circ s \circ (d_{\bar{X}, X}^2)^*.$$

Now,

$$\begin{aligned}
 [\overline{X} \otimes X, t] &= [\overline{X} \otimes X, E_1(r^*) \circ E_1(r) \circ t] = [\mathbf{1}, E_1(r) \circ t \circ E_2(r^*)] \\
 &= \left[\mathbf{1}, E_1(r^*) \circ \left(\text{id}_{E_1(\overline{X})} \otimes E_2(r^*) \circ \text{id}_{E_1(\overline{X})} \otimes s^* \otimes \text{id}_{E_2(\overline{X})} \right. \right. \\
 &\quad \left. \left. \circ E_1(r) \otimes \text{id}_{E_2(\overline{X})} \right) \otimes s \circ E_2(r) \right] \\
 &= [\mathbf{1}, E_1(r^*) \circ \text{id} \otimes (s \circ s^*) \circ E_1(r)] = [\mathbf{1}, u^*u],
 \end{aligned}$$

where we have used the conjugate equations and put $u = \text{id} \otimes s^* \circ E_1(r)$. Thus, $[X, s]^* \cdot [X, s] = [\mathbf{1}, u^*u]$ is zero iff u^*u is zero. By positivity of the $*$ -operation in \mathcal{H} , this holds iff $u = 0$. Using once again the conjugate equations we see that this is equivalent to $s = 0$. Thus for elements $a \in A(E_1, E_2)$ of the form $[X, s]$, the implication $a^*a = 0 \Rightarrow a = 0$ holds. For a general $a = \sum_i [X_i, s_i]$ we pick isometries $v_i : X_i \rightarrow X$ such that $\sum_i v_i \circ v_i^* = \text{id}_X$ (i.e. $X \cong \oplus_i X_i$). Then $[X_i, s_i] = [X, E_1(v_i) \circ s_i \circ E_2(v_i^*)]$, thus

$$\sum_i [X_i, s_i] = [X, \sum_i E_1(v_i) \circ s_i \circ E_2(v_i^*)],$$

implying that every element of $A(E_1, E_2)$ can be written as $[X, s]$, and we are done. ■

PROPOSITION 394. *Let \mathcal{C} be a TC^* and let $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{H}$ be $*$ -preserving fiber functors. Then*

$$\|a\| = \inf_b \sup_{X \in \mathcal{C}} \|b_X\|_{\text{End}E(X)},$$

where the infimum is over all representers $b \in A_0(E_1, E_2)$ of $a \in A(E_1, E_2)$, defines a C^* -norm on $A(E_1, E_2)$.

Proof. Let $[X, s], [Y, t] \in A_0(E_1, E_2)$. Then $[X, s] \cdot [Y, t] = [X \otimes Y, u]$, where $u = d_{X,Y}^1 \circ s \otimes t \circ (d_{X,Y}^2)^{-1}$. Since $d_{X,Y}^1, d_{X,Y}^2$ are unitaries, we have $\|[X \otimes Y, u]\| = \|u\| \leq \|s\| \cdot \|t\|$. Thus $\|b\| = \sup_{X \in \mathcal{C}} \|b_X\|_{\text{End}E(X)}$ defines a submultiplicative norm on $A_0(E_1, E)$, and the above formula for $\|a\|$ is the usual definition of a norm on the quotient algebra $A_0(E_1, E_2)/I(E_1, E_2)$. This norm satisfies $\|[X, s]\| = \|s\|$. Since every $a \in A(E_1, E_2)$ can be written as $[X, s]$, we have $\|a\| = 0 \Rightarrow a = 0$. Finally, the computations in the proof of Proposition 393 imply

$$\|[X, s]^*[X, s]\| = \|[\mathbf{1}, u^*u]\| = \|u^*u\| = \|u\|^2 = \|s\|^2 = \|[X, s]\|^2,$$

which is the C^* -condition. ■

DEFINITION 395. Let \mathcal{C} be a TC^* and let $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{H}$ be $*$ -preserving fiber functors. Then $\mathcal{A}(E_1, E_2)$ denotes the $\|\cdot\|$ -completion of $A(E_1, E_2)$. (This is a unital C^* -algebra, which is commutative if \mathcal{C}, E_1, E_2 are symmetric.)

B.4 Uniqueness of fiber functors

LEMMA 396. [Joyal and Street, 1993a] Let \mathcal{C} be a TC^* , \mathcal{D} a strict tensor category and $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{D}$ strict tensor functors. Then any monoidal natural transformation $\alpha : E_1 \rightarrow E_2$ is a natural isomorphism.

Proof. It is sufficient to show that every component $\alpha_X : E_1(X) \rightarrow E_2(X)$ has a two-sided inverse $\beta_X : E_2(X) \rightarrow E_1(X)$. The family $\{\beta_X, X \in \mathcal{C}\}$ will then automatically be a natural transformation. If (\bar{X}, r, \bar{r}) is a conjugate for X , monoidality of α implies

$$(58) \quad E_2(r^*) \circ \alpha_{\bar{X}} \otimes \alpha_X = E_2(r^*) \circ \alpha_{\bar{X} \otimes X} = \alpha_1 \circ E_1(r^*) = E_1(r^*).$$

If we now define

$$\beta_X = \text{id}_{E_1(X)} \otimes E_2(r^*) \circ \text{id}_{E_1(X)} \otimes \alpha_{\bar{X}} \otimes \text{id}_{E_2(X)} \circ E_1(\bar{r}) \otimes \text{id}_{E_2(X)},$$

we have

$$\begin{aligned} \beta_X \circ \alpha_X &= (\text{id}_{E_1(X)} \otimes E_2(r^*) \circ \text{id}_{E_1(X)} \otimes \alpha_{\bar{X}} \otimes \text{id}_{E_2(X)} \circ E_1(\bar{r}) \otimes \text{id}_{E_2(X)}) \circ \alpha_X \\ &= \text{id}_{E_1(X)} \otimes E_2(r^*) \circ \text{id}_{E_1(X)} \otimes \alpha_{\bar{X}} \otimes \alpha_X \circ E_1(\bar{r}) \otimes \text{id}_{E_1(X)} \\ &= \text{id}_{E_1(X)} \otimes E_1(r^*) \circ E_1(\bar{r}) \otimes \text{id}_{E_1(X)} = \text{id}_{E_1(X)}. \end{aligned}$$

The argument for $\alpha_X \circ \beta_X = \text{id}_{E_2(X)}$ is similar. ■

REMARK 397. The lemma remains correct if one allows E_1, E_2 (or even \mathcal{C}, \mathcal{D}) to be non-strict. To adapt the proof one must replace $E_1(r)$ (which is a morphism $E_1(\mathbf{1}) \rightarrow E_1(\bar{X} \otimes X)$) by $(d_{\bar{X}, X}^{E_1})^{-1} \circ E_1(r) \circ e^{E_1}$ (which is a morphism $\mathbf{1}_{\text{Vect}} \rightarrow E_1(\bar{X}) \otimes E_1(X)$). Similarly with $E_2(\bar{r})$.

PROPOSITION 398. Let \mathcal{C} be a TC^* and $E_1, E_2 : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ fiber functors. The pairing between $A_0(E_1, E_2)$ and the vector space

$$\begin{aligned} \text{Nat}(E_1, E_2) &= \left\{ (\alpha_X)_{X \in \mathcal{C}} \in \prod_{X \in \mathcal{C}} \text{Hom}(E_1(X), E_2(X)) \mid \right. \\ &\quad \left. E_2(s) \circ \alpha_X = \alpha_Y \circ E_1(s) \quad \forall s : X \rightarrow Y \right\} \end{aligned}$$

of natural transformations $E_1 \rightarrow E_2$ that is given, for $(\alpha_X) \in \text{Nat}(E_1, E_2)$ and $a \in A_0(E_1, E_2)$, by

$$(59) \quad \langle \alpha, a \rangle = \sum_{X \in \mathcal{C}} \text{Tr}_{E_1(X)}(a_X \alpha_X)$$

descends to a pairing between $\text{Nat}(E_1, E_2)$ and the quotient algebra $A(E_1, E_2) = A_0(E_1, E_2)/I(E_1, E_2)$ such that

$$\text{Nat}(E_1, E_2) \cong A(E_1, E_2)^*.$$

Under this isomorphism, an element $a \in A(E_1, E_2)^*$ corresponds to an element of $\text{Nat}_{\otimes}(E_1, E_2)$, i.e. a monoidal natural transformation (thus isomorphism by Lemma 396), iff it is a character, to wit multiplicative.

Proof. The dual vector space of the direct sum $A_0(E_1, E_2)$ is the direct product $\prod_{X \in \mathcal{C}} \text{Hom}(E_2(X), E_1(X))^*$, and since the pairing between $\text{Hom}(E_2(X), E_1(X)) \times \text{Hom}(E_1(X), E_2(X))$, $s \times t \mapsto \text{Tr}(s \circ t)$ is non-degenerate, we have

$$A_0(E_1, E_2)^* \cong \prod_{X \in \mathcal{C}} \text{Hom}(E_1(X), E_2(X))$$

w.r.t. the pairing given in (59). Now, $A(E_1, E_2)$ is the quotient of $A_0(E_1, E_2)$ by the subspace $I(E_1, E_2)$, thus the dual space $A(E_1, E_2)^*$ consists precisely of those elements of $A_0(E_1, E_2)^*$ that are identically zero on $I(E_1, E_2)$. Assume $(\alpha_X)_{X \in \mathcal{C}}$ satisfies $\langle \alpha, a \rangle = 0$ for all $a \in I(E_1, E_2)$, equivalently $\langle \alpha, [X, a \circ E_2(s)] - [Y, E_1(s) \circ a] \rangle = 0$ for all $s : X \rightarrow Y$ and $a : E_2(Y) \rightarrow E_1(X)$. By definition (59) of the pairing, this is equivalent to

$$\text{Tr}_{E_1(X)}(a \circ E_2(s) \circ \alpha_X) - \text{Tr}_{E_1(Y)}(E_1(s) \circ a \circ \alpha_Y) = 0$$

for all $s : X \rightarrow Y$ and $a \in \text{Hom}(E_2(Y), E_1(X))$. Non-degeneracy of the trace implies that $\alpha = (\alpha_X)_{X \in \mathcal{C}}$ must satisfy $E_2(s) \circ \alpha_X = \alpha_Y \circ E_1(s)$ for all $s : X \rightarrow Y$, thus $\alpha \in \text{Nat}(E_1, E_2)$, implying

$$A(E_1, E_2)^* \cong \text{Nat}(E_1, E_2).$$

Now we consider the question when the functional $\phi \in A(E_1, E_2)^*$ corresponding to $\alpha \in \text{Nat}(E_1, E_2)$ is a character, i.e. multiplicative. This is the case when

$$\langle \alpha, [X, s] \cdot [Y, t] \rangle = \langle \alpha, [X, s] \rangle \langle \alpha, [Y, t] \rangle \quad \forall [X, s], [Y, t] \in A(E_1, E_2).$$

(Strictly speaking, $[X, s], [Y, t]$ are representers in $A_0(E_1, E_2)$ for some elements in $A(E_1, E_2)$.) In view of (59) and the definition of the product in $A(E_1, E_2)$ this amounts to

$$\begin{aligned} \text{Tr}_{E_1(X \otimes Y)}(d_{X,Y}^1 \circ s \otimes t \circ (d_{X,Y}^2)^{-1} \circ \alpha_{X \otimes Y}) &= \text{Tr}_{E_1(X)}(s \circ \alpha_X) \text{Tr}_{E_1(Y)}(t \circ \alpha_Y) \\ &= \text{Tr}_{E_1(X) \otimes E_2(X)}((s \circ \alpha_X) \otimes (t \circ \alpha_Y)) \\ &= \text{Tr}_{E_1(X) \otimes E_2(X)}(s \otimes t \circ \alpha_X \otimes \alpha_Y). \end{aligned}$$

In view of the cyclic invariance and non-degeneracy of the trace, this is true for all $s : E_2(X) \rightarrow E_1(X)$ and $t : E_2(Y) \rightarrow E_1(Y)$, iff

$$\alpha_{X \otimes Y} = d_{X,Y}^2 \circ \alpha_X \otimes \alpha_Y \circ (d_{X,Y}^1)^{-1} \quad \forall X, Y \in \mathcal{C}.$$

This is precisely the condition for $\alpha \in \text{Nat}(E_1, E_2)$ to be monoidal, to wit $\alpha \in \text{Nat}_{\otimes}(E_1, E_2)$. ■

PROPOSITION 399. *Let \mathcal{C} be a TC^* and let $E_1, E_2 : \mathcal{C} \rightarrow \mathcal{H}$ be $*$ -preserving fiber functors. Then a monoidal natural transformation $\alpha \in \text{Nat}_{\otimes}(E_1, E_2)$ is unitary (i.e. each α_X is unitary) iff the corresponding character $\phi \in A(E_1, E_2)$ is a $*$ -homomorphism (i.e. $\phi(a^*) = \overline{\phi(a)}$).*

Proof. Let $\alpha \in \text{Nat}_\otimes(E_1, E_2)$ and $[X, s] \in A(E_1, E_2)$. By definition of the pairing of $A(E_1, E_2)$ and $\text{Nat}(E_1, E_2)$,

$$\phi([X, s]) = \langle \alpha, [X, s] \rangle = \text{Tr}_{E_1(X)}(s \circ \alpha_X),$$

and therefore, using $\overline{\text{Tr}(AB)} = \text{Tr}(A^*B^*)$,

$$\overline{\phi([X, s])} = \text{Tr}_{E_1(X)}(s^* \circ \alpha_X^*).$$

On the other hand,

$$\begin{aligned} \phi([X, s]^*) &= \langle \alpha, [\overline{X}, \text{id}_{E_1(\overline{X})} \otimes E_2(\overline{r}^*) \circ \text{id}_{E_1(\overline{X})} \otimes s^* \otimes \text{id}_{E_2(\overline{X})} \circ E_1(r) \otimes \text{id}_{E_2(\overline{X})}] \rangle \\ &= \text{Tr}_{E_1(\overline{X})}(\text{id}_{E_1(\overline{X})} \otimes E_2(\overline{r}^*) \circ \text{id}_{E_1(\overline{X})} \otimes s^* \otimes \text{id}_{E_2(\overline{X})} \circ E_1(r) \otimes \text{id}_{E_2(\overline{X})} \circ \alpha_{\overline{X}}) \\ &= E_2(\overline{r}^*) \circ s^* \otimes \alpha_{\overline{X}} \circ E_1(\overline{r}) \\ &= E_2(\overline{r}^*) \circ (\alpha_X \circ \alpha_X^{-1} \circ s^*) \otimes \alpha_{\overline{X}} \circ E_1(\overline{r}) \\ &= E_1(\overline{r}^*) \circ (\alpha_X^{-1} \circ s^*) \otimes \text{id}_{E_2(\overline{X})} \circ E_1(\overline{r}) \\ &= \text{Tr}_{E_1(X)}(\alpha_X^{-1} \circ s^*). \end{aligned}$$

(In the fourth step we have used the invertibility of α (Lemma 396) and in the fifth equality we have used (58) with X and \overline{X} interchanged and r replaced by \overline{r} .) Now non-degeneracy of the trace implies that $\overline{\phi([X, s])} = \phi([X, s]^*)$ holds for all $[X, s] \in (E_1, E_2)$ iff $\alpha_X^* = \alpha_{\overline{X}}^{-1}$ for all $X \in \mathcal{C}$, as claimed. ■

Now we are in a position to prove the first of our outstanding claims:

Proof of Theorem 373. By the preceding constructions, the $\|\cdot\|$ -closure $\mathcal{A}(E_1, E_2)$ of $A(E_1, E_2)$ is a commutative unital C^* -algebra. As such it has (lots of) characters, i.e. unital $*$ -homomorphisms into \mathbb{C} . (Cf. e.g. Theorem 401 below.) Such a character restricts to $A(E_1, E_2)$ and corresponds, by Propositions 398 and 399, to a unitary monoidal natural transformation $\alpha \in \text{Nat}(E_1, E_2)$. ■

REMARK 400. 1. The discussion of the algebra $A(E_1, E_2)$ is inspired by the one in the preprint [Bichon, ND] that didn't make it into the published version [Bichon, 1998]. The above proof of Theorem 373 first appeared in [Bichon, 1999].

2. Lemma 396 implies that the category consisting of fiber functors and monoidal natural transformations is a groupoid, i.e. every morphism is invertible. Theorem 373 then means that the category consisting of symmetric $*$ -preserving fiber functors and unitary monoidal natural transformations is a transitive groupoid, i.e. all objects are isomorphic. That this groupoid is non-trivial is the statement of Theorem 382, whose proof will occupy the bulk of this section, beginning in Subsection B.6.

B.5 The concrete Tannaka theorem. Part II

In order to prove Proposition 375 we need the formalism of the preceding subsections. We write $\mathcal{A}(E)$ for the commutative unital C^* -algebra $\mathcal{A}(E, E)$ defined

earlier. In order to study this algebra we need some results concerning commutative unital C^* -algebras that can be gathered, e.g., from [Pedersen, 1989].

THEOREM 401. *Let \mathcal{A} be a commutative unital C^* -algebra. Let \mathcal{A}^* be its Banach space dual and let*

$$\begin{aligned} P(\mathcal{A}) &= \{\phi \in \mathcal{A}^* \mid \phi(1) = 1, \|\phi\| \leq 1\}, \\ X(\mathcal{A}) &= \{\phi \in \mathcal{A}^* \mid \phi(1) = 1, \phi(ab) = \phi(a)\phi(b), \phi(a^*) = \overline{\phi(a)} \ \forall a, b \in \mathcal{A}\}. \end{aligned}$$

$P(\mathcal{A})$ and $X(\mathcal{A})$ are equipped with the w^* -topology on \mathcal{A} according to which $\phi_\iota \rightarrow \phi$ iff $\phi_\iota(a) \rightarrow \phi(a)$ for all $a \in \mathcal{A}$. Then:

- (i) $X(\mathcal{A}) \subset P(\mathcal{A})$. (I.e., $*$ -characters have norm ≤ 1 .)
- (ii) $X(\mathcal{A})$ is compact w.r.t. the w^* -topology on $P(\mathcal{A})$.
- (iii) The map $\mathcal{A} \rightarrow C(X(\mathcal{A}))$ given by $a \mapsto (\phi \mapsto \phi(a))$ is an isomorphism of C^* -algebras.
- (iv) The convex hull

$$\left\{ \sum_{i=1}^N c_i \phi_i, N \in \mathbb{N}, c_i \in \mathbb{R}_+, \sum_i c_i = 1, \phi_i \in X(\mathcal{A}) \right\}$$

of $X(\mathcal{A})$ is w^* -dense in $P(\mathcal{A})$.

Proof. (i) Any unital $*$ -homomorphism α of Banach algebras satisfies $\|\alpha(a)\| \leq \|a\|$.

(ii) By Alaoglu’s theorem [Pedersen, 1989, Theorem 2.5.2], the unit ball of \mathcal{A}^* is compact w.r.t. the w^* -topology, and so are the closed subsets $X(\mathcal{A}) \subset P(\mathcal{A}) \subset \mathcal{A}^*$.

(iii) This is Gelfand’s theorem, cf. [Pedersen, 1989, Theorem 4.3.13].

(iv) This is the Krein-Milman theorem, cf. Theorem 2.5.4 together with Proposition 2.5.7 in [Pedersen, 1989]. ■

Theorem 401, (ii) implies that the set $X \equiv X(\mathcal{A}(E))$ of $*$ -characters of $\mathcal{A}(E)$ is a compact Hausdorff space w.r.t. the w^* -topology. By (iii) and Proposition 399, the elements of X are in bijective correspondence with the set G_E of unitary monoidal transformations of the functor E .

LEMMA 402. *The bijection $X \cong G_E$ is a homeomorphism w.r.t. the topologies defined above.*

Proof. By definition of the product topology on $\prod_{X \in \mathcal{C}} \mathcal{U}(E(X))$, a net (g_ι) in G_E converges iff the net $(g_{\iota, X})$ in $\mathcal{U}(E(X))$ converges for every $X \in \mathcal{C}$. On the other hand, a net (ϕ_ι) in X converges iff $(\phi_\iota(a))$ converges in \mathbb{C} for every $a \in \mathcal{A}(E)$. In view of the form of the correspondence $\phi \leftrightarrow g$ established in Proposition 398, these two notions of convergence coincide. ■

The homeomorphism $X \cong G_E$ allows to transfer the topological group structure that G_E automatically has to the compact space X . Now we are in a position to complete the proof of our second outstanding claim.

Proof of Proposition 375. Since \mathcal{C} is semisimple and essentially small, there exist a set I and a family $\{X_i, i \in I\}$ of irreducible objects such that every object is (isomorphic to) a finite direct sum of objects from this set. If $\text{Nat}(E) \equiv \text{Nat}(E, E)$ is the space of natural transformations from E to itself, with every $\alpha \in \text{Nat}(E)$ we can associate the family $(\alpha_i = \alpha_{X_i})_{i \in I}$, which is an element of $\prod_{i \in I} \text{End } E(X_i)$. Semisimplicity of \mathcal{C} and naturality of α imply that every such element arises from exactly one natural transformation of E . (In case it is not obvious, a proof can be found in [Müger *et al.*, 2004, Proposition 5.4].) In this way we obtain an isomorphism

$$\gamma : \text{Nat}(E) \rightarrow \prod_{i \in I} \text{End } E(X_i), \quad \alpha \mapsto (\alpha_{X_i})_{i \in I}$$

of vector spaces. Now consider the linear map

$$\delta : \bigoplus_{i \in I} \text{End } E(X_i) \rightarrow A(E), \quad (a_i) \mapsto \sum_i [X_i, a_i].$$

Since every $a \in A(E)$ can be written as $[X, s]$ (proof of Proposition 393) and every $[X, s]$ is a sum of elements $[X_i, s_i]$ with X_i irreducible, δ is surjective. When understood as a map to $A_0(E)$, δ obviously is injective. As a consequence of $\text{Hom}(X_i, X_j) = \{0\}$ for $i \neq j$, the image in $A_0(E)$ of δ has trivial intersection with the ideal $I(E)$, which is the kernel of the quotient map $A_0(E) \rightarrow A(E)$, thus δ is injective and therefore an isomorphism (of vector spaces, not algebras). If the C^* -norm on $A(E)$ is pulled back via δ we obtain the norm

$$\|(a_i)_{i \in I}\| = \sup_{i \in I} \|a_i\|_{\text{End } E(X_i)}$$

on $\bigoplus_{i \in I} \text{End } E(X_i)$. Thus we have an isomorphism $\bar{\delta} : \overline{\bigoplus_{i \in I} \text{End } E(X_i)}^{\|\cdot\|} \rightarrow \mathcal{A}(E)$ of the norm closures. W.r.t. the isomorphisms γ, δ , the pairing $\langle \cdot, \cdot \rangle : \text{Nat}(E) \times A(E) \rightarrow \mathbb{C}$ of Proposition 398 becomes

$$\langle \cdot, \cdot \rangle^\sim : \prod_{i \in I} \text{End } E(X_i) \times \bigoplus_{i \in I} \text{End } E(X_i) \rightarrow \mathbb{C}, \quad (\alpha_{X_i}) \times (a_i) \mapsto \sum_{i \in I} \text{Tr}_{E(X_i)}(\alpha_i a_i).$$

(More precisely: $\langle \cdot, \delta(\cdot) \rangle = \langle \gamma(\cdot), \cdot \rangle^\sim$ as maps $\text{Nat}(E) \times \bigoplus_{i \in I} \text{End } E(X_i) \rightarrow \mathbb{C}$.) Thus if $\alpha \in \text{Nat}(E)$ is such that $\gamma(\alpha) \in \prod_{i \in I} \text{End } E(X_i)$ has only finitely many non-zero components (i.e. $\gamma(\alpha) \in \bigoplus_{i \in I} \text{End } E(X_i)$), then $\langle \alpha, \cdot \rangle \in A(E)^*$ extends to an element of $\mathcal{A}(E)^*$.

Now (iv) of Theorem 401 implies that every $\phi \in \mathcal{A}(E)^*$ is the w^* -limit of a net (ϕ_ι) in the \mathbb{C} -span of the $*$ -characters $X(\mathcal{A}(E))$ of $\mathcal{A}(E)$. Thus for every $(\alpha_i) \in \bigoplus_{i \in I} \text{End } E(X_i)$ there is a such a net (ϕ_ι) for which

$$w^* - \lim \phi_\iota = \langle \gamma^{-1}((\alpha_i)), \cdot \rangle \in \mathcal{A}(E)^*.$$

Restricting the ϕ_i to $A(E)$ and using the isomorphism $\text{Nat } E \cong A(E)^*$, we obtain a net in $\text{Nat } E$ that converges to $\gamma^{-1}((\alpha_i))$. By Propositions 398, 399, the isomorphism $A(E)^* \rightarrow \text{Nat } E$ maps the elements of $X(\mathcal{A}(E))$ to the unitary natural monoidal transformations of E , i.e. to elements of G_E . Thus, in particular for every finite $S \subset I$ we have

$$\overline{\text{span}}_{\mathbb{C}} \left\{ \underbrace{\pi_{s_1}(g) \oplus \cdots \oplus \pi_{s_{|S|}}(g)}_{\text{all } s \in S}, g \in G_E \right\} = \bigoplus_{s \in S} \text{End } E(X_s),$$

which clearly is a good deal more than claimed in Proposition 375. ■

This concludes the proof of all ingredients that went into the proof of Theorem 377. From the proof it is obvious that the commutative C^* -algebra $\mathcal{A}(E)$ is just the algebra of continuous functions on the compact group G_E , whereas $A(E)$ is the linear span of the matrix elements of the finite dimensional representations of G_E .

B.6 Making a symmetric fiber functor $*$ -preserving

The aim of his subsection is to prove the following result, which seems to be new:

THEOREM 403. *An even STC* \mathcal{C} that admits a symmetric fiber functor $\mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ also admits a symmetric $*$ -preserving fiber functor $\mathcal{C} \rightarrow \mathcal{H}$.*

LEMMA 404. *Let \mathcal{C} be an STC* and $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ a symmetric fiber functor. Choose arbitrary positive definite inner products $\langle \cdot, \cdot \rangle_X^0$ (i.e. Hilbert space structures) on all of the spaces $E(X), X \in \mathcal{C}$. Then the maps $X \mapsto E(X)$ and $s \mapsto E(s^*)^\dagger$, where $E(s^*)^\dagger$ is the adjoint of $E(s^*)$ w.r.t. the inner products $\langle \cdot, \cdot \rangle_X^0$, define a faithful functor $\tilde{E} : \mathbb{C} \rightarrow \text{Vect}_{\mathbb{C}}$. With $d_{X,Y}^{\tilde{E}} = ((d_{X,Y}^E)^\dagger)^{-1}$ and $e^{\tilde{E}} = ((e^E)^\dagger)^{-1}$, this is a symmetric fiber functor.*

Proof. First note that $s \mapsto \tilde{E}(s)$ is \mathbb{C} -linear and really defines a functor, since $\tilde{E}(\text{id}_X) = \text{id}_{\tilde{E}(X)}$ and

$$\tilde{E}(s \circ t) = E((s \circ t)^*)^\dagger = E(t^* \circ s^*)^\dagger = (E(t^*) \circ E(s^*))^\dagger = E(s^*)^\dagger \circ E(t^*)^\dagger = \tilde{E}(s) \circ \tilde{E}(t).$$

Faithfulness of E clearly implies faithfulness of \tilde{E} . With $d_{X,Y}^{\tilde{E}} = ((d_{X,Y}^E)^\dagger)^{-1}$ and $e^{\tilde{E}} = ((e^E)^\dagger)^{-1}$, commutativity of the diagrams (56) and (57) is obvious. Since E is a tensor functor, we have

$$E(s \otimes t) \circ d_{X,Y}^E = d_{X',Y'}^E \circ E(s) \otimes E(t)$$

for all $s : X \rightarrow X', t : Y \rightarrow Y'$, which is equivalent to

$$(E(s \otimes t))^\dagger \circ ((d_{X',Y'}^E)^{-1})^\dagger = ((d_{X,Y}^E)^{-1})^\dagger \circ (E(s) \otimes E(t))^\dagger.$$

Since this holds for all s, t , we have proven naturality of the family $(d_{X,Y}^{\tilde{E}})$, thus \tilde{E} is a tensor functor. The computation

$$\tilde{E}(c_{X,Y}) = E(c_{X,Y}^*)^\dagger = E(c_{Y,X})^\dagger = \Sigma_{E(Y),E(X)}^\dagger = \Sigma_{E(X),E(Y)},$$

where we have used $\Sigma_{H,H'}^\dagger = \Sigma_{H',H}$, shows that \tilde{E} is also symmetric. Thus \tilde{E} is a symmetric fiber functor. ■

Now the discussion of Subsection B.3 applies and provides us with a commutative unital \mathbb{C} -algebra $A(E, \tilde{E})$. However, we cannot appeal to Proposition 393 to conclude that $A(E, \tilde{E})$ is a \star -algebra, since E, \tilde{E} are not \star -preserving. In fact, for arbitrary symmetric fiber functors E_1, E_2 there is no reason for the existence of a positive \star -operation on $A(E_1, E_2)$, but in the present case, where the two functors are related by $E_2(s) = E_1(s^*)^\dagger$, this is true:

PROPOSITION 405. *Let \mathcal{C} be an STC^* , $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ a symmetric fiber functor and \tilde{E} as defined above. Then*

$$[X, s]^\star = [X, s^\dagger]$$

is well defined and is a positive \star -operation on $A(E, \tilde{E})$. With respect to this \star -operation, the norm $\| \cdot \|$ from Proposition 394 is a C^ -norm, i.e. $\|a^\star a\| = \|a\|^2$ for all $a \in A(E, \tilde{E})$.*

Proof. For $[X, s] \in A_0(E, \tilde{E})$ we define $[X, s]^\star = [X, s^\dagger]$, where s^\dagger is the adjoint of $s \in \text{End}E(X)$ w.r.t. the inner product on $E(X)$. Clearly, \star is involutive and antilinear. Now, if $s : X \rightarrow Y$, $a \in \text{Hom}(E_2(Y), E_1(X))$, then

$$\begin{aligned} ([X, a \circ E_2(s)] - [Y, E_1(s) \circ a])^\star &= [X, a \circ E(s^*)^\dagger]^\star - [Y, E(s) \circ a]^\star \\ &= [X, E(s^*) \circ a^\dagger] - [Y, a^\dagger \circ E(s)^\dagger] = [X, E_1(s^*) \circ a^\dagger] - [Y, a^\dagger \circ E_2(s^*)]. \end{aligned}$$

Since $s^* \in \text{Hom}(Y, X)$ and $a^\dagger \in \text{Hom}(E(X), E(Y))$, the right hand side of this expression is again in $I(E, \tilde{E})$. Thus $I(E, \tilde{E})$ is stable under \star , and \star descends to an antilinear involution on $A(E, \tilde{E})$. In $A_0(E, \tilde{E})$ we have

$$\begin{aligned} ([X, s] \cdot [Y, t])^\star &= [X \otimes Y, d_{X,Y}^{\tilde{E}} \circ s \otimes t \circ (d_{X,Y}^E)^{-1}]^\star \\ &= [X \otimes Y, (d_{X,Y}^E)^\dagger]^{-1} \circ s \otimes t \circ (d_{X,Y}^E)^{-1}]^\star \\ &= [X \otimes Y, (d_{X,Y}^E)^\dagger]^{-1} \circ s^\dagger \otimes t^\dagger \circ (d_{X,Y}^E)^{-1}] \\ &= [X \otimes Y, d_{X,Y}^{\tilde{E}} \circ s^\dagger \otimes t^\dagger \circ (d_{X,Y}^E)^{-1}] \\ &= [X, s]^\star \cdot [Y, t]^\star. \end{aligned}$$

Together with commutativity of $A(E, \tilde{E})$ this implies that \star is antimultiplicative. Recall that there is an isomorphism $\delta : \bigoplus_{i \in I} \text{End} E(X_i) \rightarrow A(E, \tilde{E})$ such that $\|\delta((a_i)_{i \in I})\| = \sup_i \|a_i\|$, where $\| \cdot \|$ is the norm defined in Subsection B.3. By definition of \star we have $\delta((a_i)^\star) = \delta((a_i^\dagger))$, implying $\|a^\star a\| = \|a\|^2$. Thus $(A(E, \tilde{E}), \star, \| \cdot \|)$ is a pre- C^* -algebra. ■

(Note that the involution \star has nothing at all to do with the one defined in Subsection B.3!)

PROPOSITION 406. *Let \mathcal{C} be an STC^* and $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ a symmetric fiber functor. With \tilde{E} as defined above, there exists a natural monoidal isomorphism $\alpha : E \rightarrow \tilde{E}$, whose components α_X are positive, i.e. $\langle u, \alpha_X u \rangle_X^0 > 0$ for all nonzero $u \in E(X)$.*

Proof. As in Subsection B.4, the norm-completion $\mathcal{A}(E, \tilde{E})$ of $A(E, \tilde{E})$ is a commutative unital C^* -algebra and therefore admits a $*$ -character $\phi : \mathcal{A}(E, \tilde{E}) \rightarrow \mathbb{C}$. Restricting to $A(E, \tilde{E})$, Proposition 398 provides a monoidal natural isomorphism $\alpha : E \rightarrow \tilde{E}$. But we know more: The character ϕ is positive, i.e. $\phi(a^*a) > 0$ for all $a \neq 0$. With $a = [X, s]$ and taking (59) into account, we have

$$\begin{aligned} \phi(a^*a) &= \phi([X, s^\dagger s]) = \text{Tr}_{E(X)}(s^\dagger s \alpha_X) = \text{Tr}_{E(X)}(s \alpha_X s^\dagger) \\ &= \sum_i \langle e_i, s \alpha_X s^\dagger e_i \rangle_X^0 = \sum_i \langle s^\dagger e_i, \alpha_X s^\dagger e_i \rangle_X^0, \end{aligned}$$

where $\{e_i\}$ is any basis of $E(X)$ that is orthonormal w.r.t. $\langle \cdot, \cdot \rangle_X^0$. This is positive for all $a = [X, s] \in \mathcal{A}(E, \tilde{E})$ iff $\langle u, \alpha_X u \rangle_X^0 > 0$ for all nonzero $u \in E(X)$. ■

Now we are in a position to prove the main result of this subsection, which is a more specific version of Theorem 403.

THEOREM 407. *Let \mathcal{C} be an even STC^* and $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ a symmetric fiber functor. Then there exist Hilbert space structures (i.e. positive definite inner products $\langle \cdot, \cdot \rangle_X$) on the spaces $E(X), X \in \mathcal{C}$ such that $X \mapsto (E(X), \langle \cdot, \cdot \rangle_X)$ is a $*$ -preserving symmetric fiber functor $\mathcal{C} \rightarrow \mathcal{H}$.*

Proof. Pick non-degenerate inner products $\langle \cdot, \cdot \rangle_X^0$ on the spaces $E(X), X \in \mathcal{C}$. Since $E(\mathbf{1})$ is one-dimensional and spanned by $e^E \mathbf{1}$, where $1 \in \mathbb{C} = \mathbf{1}_{\text{Vect}_{\mathbb{C}}}$, we can define $\langle \cdot, \cdot \rangle_{\mathbf{1}}^0$ by $\langle ae^E \mathbf{1}, be^E \mathbf{1} \rangle_{\mathbf{1}}^0 = \bar{a}b$, as will be assumed in the sequel. Let \tilde{E} and $\alpha \in \text{Nat}_{\otimes}(E, \tilde{E})$ as above. Defining new inner products $\langle \cdot, \cdot \rangle_X$ on the spaces $E(X)$ by

$$\langle v, u \rangle_X = \langle v, \alpha_X u \rangle_X^0,$$

the naturality

$$\alpha_Y \circ E(s) = \tilde{E}(s) \circ \alpha_X = E(s^*)^\dagger \circ \alpha_X \quad \forall s : X \rightarrow Y$$

of (α_X) implies

$$\begin{aligned} \langle v, E(s)u \rangle_Y &= \langle v, \alpha_Y E(s)u \rangle_Y^0 = \langle v, E(s^*)^\dagger \alpha_X u \rangle_Y^0 \\ &= \langle E(s^*)v, \alpha_X u \rangle_X^0 = \langle E(s^*)v, u \rangle_X \end{aligned}$$

for all $s : X \rightarrow Y, u \in E(X), v \in E(Y)$. This is the same as $E(s^*) = E(s)^*$, where now $E(s)^*$ denotes the adjoint of $E(s)$ w.r.t. the inner products $\langle \cdot, \cdot \rangle$.

Thus the functor $X \mapsto (E(X), \langle \cdot, \cdot \rangle_X)$ is $*$ -preserving. The new inner products $\langle \cdot, \cdot \rangle_X$ are non-degenerate since the α_X are invertible, and the positivity property $\langle u, \alpha_X u \rangle_X^0 > 0$ for $u \neq 0$ implies that $(E(X), \langle \cdot, \cdot \rangle_X)$ is a Hilbert space. The monoidality

$$\alpha_{X \otimes Y} \circ d_{X,Y}^E = d_{X,Y}^{\tilde{E}} \circ \alpha_X \otimes \alpha_Y = ((d_{X,Y}^E)^\dagger)^{-1} \circ \alpha_X \otimes \alpha_Y \quad \forall X, Y$$

of the natural isomorphism $\alpha : E \rightarrow \tilde{E}$ is equivalent to

$$(60) \quad \alpha_X \otimes \alpha_Y = (d_{X,Y}^E)^\dagger \circ \alpha_{X \otimes Y} \circ d_{X,Y}^E.$$

Using this we have

$$\begin{aligned} \langle d_{X,Y}^E(u' \otimes v'), d_{X,Y}^E(u \otimes v) \rangle_{X \otimes Y} &= \langle d_{X,Y}^E(u' \otimes v'), \alpha_{X \otimes Y} \circ d_{X,Y}^E(u \otimes v) \rangle_{X \otimes Y}^0 \\ &= \langle (u' \otimes v'), (d_{X,Y}^E)^\dagger \circ \alpha_{X \otimes Y} \circ d_{X,Y}^E(u \otimes v) \rangle_{X \otimes Y}^0 \\ &= \langle (u' \otimes v'), (\alpha_X \otimes \alpha_Y)(u \otimes v) \rangle_{X \otimes Y}^0 \\ &= \langle u', \alpha_X u \rangle_X^0 \langle v', \alpha_Y v \rangle_Y^0 = \langle u', u \rangle_X \langle v', v \rangle_Y, \end{aligned}$$

thus the isomorphisms $d_{X,Y}^E : E(X) \otimes E(Y) \rightarrow E(X \otimes Y)$ are unitary w.r.t. the inner products $\langle \cdot, \cdot \rangle$.

Now, the compatibility (57) of d^E and e^E implies that $d_{\mathbf{1},\mathbf{1}}^E \circ e^E \mathbf{1} \otimes e^E \mathbf{1} = e^E \mathbf{1}$ and therefore, using our choice of the inner product $\langle \cdot, \cdot \rangle_{\mathbf{1}}^0$,

$$\begin{aligned} \langle d_{\mathbf{1},\mathbf{1}}^E(ae^E \mathbf{1} \otimes be^E \mathbf{1}), d_{\mathbf{1},\mathbf{1}}^E(ce^E \mathbf{1} \otimes de^E \mathbf{1}) \rangle_{\mathbf{1} \otimes \mathbf{1}}^0 &= \langle abe^E \mathbf{1}, cde^E \mathbf{1} \rangle_{\mathbf{1}}^0 \\ &= \overline{abcd} = \langle ae^E \mathbf{1}, ce^E \mathbf{1} \rangle_{\mathbf{1}}^0 \langle be^E \mathbf{1}, de^E \mathbf{1} \rangle_{\mathbf{1}}^0. \end{aligned}$$

This means that $d_{\mathbf{1},\mathbf{1}}^E : E(\mathbf{1}) \otimes E(\mathbf{1}) \rightarrow E(\mathbf{1})$ is unitary w.r.t. the inner product $\langle \cdot, \cdot \rangle_{\mathbf{1}}^0$. Taking $X = Y = \mathbf{1}$ in (60) and using $\alpha_{\mathbf{1}} = \text{id}_{E(\mathbf{1})}$, we get $\lambda^2 = \lambda$. Since $\alpha_{\mathbf{1}}$ is invertible, we have $\lambda = 1$, thus $\alpha_{\mathbf{1}} = \text{id}_{E(\mathbf{1})}$ and therefore $\langle \cdot, \cdot \rangle_{\mathbf{1}} = \langle \cdot, \cdot \rangle_{\mathbf{1}}^0$. Now,

$$\langle e^E \mathbf{1}, e^E \mathbf{1} \rangle_{\mathbf{1}} = \langle e^E \mathbf{1}, \alpha_{\mathbf{1}} e^E u \rangle_{\mathbf{1}}^0 = \langle e^E \mathbf{1}, e^E \mathbf{1} \rangle_{\mathbf{1}}^0 = 1 = \langle \mathbf{1}, \mathbf{1} \rangle_{\mathbb{C}},$$

thus $(e^E)^* e^E = \text{id}_{\mathbb{C}}$. By one-dimensionality of the spaces involved, we also have $e^E (e^E)^* = \text{id}_{E(\mathbf{1})}$, thus $e^E : \mathbf{1} \rightarrow E(\mathbf{1})$ is unitary w.r.t. the inner new products $\langle \cdot, \cdot \rangle$. ■

B.7 Reduction to finitely generated categories

DEFINITION 408. An additive tensor category \mathcal{C} is finitely generated if there exists an object $Z \in \mathcal{C}$ such that every object $X \in \mathcal{C}$ is a direct summand of some tensor power $Z^{\otimes n} = \underbrace{Z \otimes \cdots \otimes Z}_{n \text{ factors}}$, $n \in \mathbb{N}$, of Z .

LEMMA 409. Let \mathcal{C} be a TC^* . Then the finitely generated tensor subcategories of \mathcal{C} form a directed system, and \mathcal{C} is the inductive limit of the latter:

$$\mathcal{C} \cong \varinjlim_{i \in I} \mathcal{C}_i.$$

Proof. Consider all full tensor subcategories of \mathcal{C} . Since \mathcal{C} is essentially small, the equivalence classes of such subcategories form a set, partially ordered by inclusion. If $\mathcal{C}_1, \mathcal{C}_2 \subset \mathcal{C}$ are finitely generated, say by the objects X_1, X_2 , then the smallest tensor subcategory containing \mathcal{C}_1 and \mathcal{C}_2 is generated by $X_1 \oplus X_2$, thus we have a directed system. Clearly there is a full and faithful tensor functor $\lim_{\leftarrow i \in I} \mathcal{C}_i \rightarrow \mathcal{C}$. Since every object X is contained in a finitely generated tensor subcategory (e.g., the one generated by X), this functor is essentially surjective and thus an equivalence of categories, cf. [Mac Lane, 1998], in fact of tensor categories, cf. [Saavedra Rivano, 1972]. ■

REMARK 410. 1. The reason for considering finitely generated categories is that the existence problem of fiber functors for such categories can be approached using powerful purely algebraic methods. The general case can then be reduced to the finitely generated one using Lemma 409.

2. Note that we don't require the generator Z to be irreducible. Thus if we a priori only know that \mathcal{C} is generated by a finite set Z_1, \dots, Z_r of objects, the direct sum $Z = \oplus_i Z_i$ will be a (reducible) generator of \mathcal{C} . This is why only a single generating object appears in the definition.

3. If G is a compact group, the category $\text{Rep}_f G$ is finitely generated iff G is a Lie group. (Proof: \Leftarrow is a consequence of the well known representation theory of compact Lie groups. \Rightarrow : It is well known that the finite dimensional representations of G separate the elements of G . Therefore, if (H, π) is a generator of $\text{Rep}_f G$, it is clear that π must be faithful. Thus G is isomorphic to a closed subgroup of the compact Lie group $\mathcal{U}(H)$, and as such it is a Lie group.)

4. The index set I in Lemma 409 can be taken countable iff \mathcal{C} has countably many isomorphism classes of irreducible objects. The category $\text{Rep}_f G$, where G is a compact group, has this property iff G is second countable, equivalently metrizable.

In Subsections B.8-B.11 we will prove the following result, which we take for granted for the moment:

THEOREM 411. *A finitely generated even STC^* admits a symmetric fiber functor $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$.*

Proof of Theorem 382. By Lemma 409, we can represent \mathcal{C} as an inductive limit $\lim_{\leftarrow i \in I} \mathcal{C}_i$ of finitely generated categories. Now Theorem 411 provides us with symmetric fiber functors $E_i : \mathcal{C}_i \rightarrow \text{Vect}_{\mathbb{C}}$, $i \in I$, and Theorem 407 turns the latter into $*$ -preserving symmetric fiber functors $E_i : \mathcal{C}_i \rightarrow \mathcal{H}$. By Theorem 377, we obtain compact groups $G_i = \text{Nat}_{\otimes} E_i$ (in fact compact Lie groups by Remark 410.3) with representations $\pi_{i,X}$ on the spaces $E_i(X)$, $X \in \mathcal{C}_i$ such that the functors $F_i : \mathcal{C}_i \rightarrow \text{Rep}_f G_i$, $X \mapsto (E_i(X), \pi_{i,X})$ are equivalences. Let now $i \leq j$, implying that \mathcal{C}_i is a full subcategory of \mathcal{C}_j . Then $E_j \upharpoonright \mathcal{C}_i$ is a fiber functor for \mathcal{C}_i and thus Theorem 373 implies the existence of a unitary natural isomorphism $\alpha^{i,j} : F_1 \rightarrow F_2 \upharpoonright \mathcal{C}_i$. (Note that $\alpha^{i,j}$ is not unique!) Now, by definition every $g \in G_2$ is a family

of unitaries $(g_X \in \mathcal{U}(E_2(X)))_{X \in \mathcal{C}_2}$ defining a monoidal natural automorphism of E_2 . Defining, for every $X \in \mathcal{C}_1$, $h_X := \alpha_X^{i,j} \circ g_X \circ (\alpha_X^{i,j})^*$ we see that the family $(h_X \in \mathcal{U}(E_1(X)))_{X \in \mathcal{C}_1}$ is a unitary monoidal natural automorphism of E_1 , to wit an element of G_1 . In this way we obtain a map $\beta^{i,j} : G_j \rightarrow G_i$ that clearly is a group homomorphism and continuous. By Schur's lemma, the unitary $\alpha_X^{i,j}$ is unique up to a phase for irreducible X . Thus for such X , $\beta_X^{i,j}$ is independent of the chosen $\alpha^{i,j}$, and thus $\beta^{i,j}$ is uniquely determined. It is also surjective in view of the Galois correspondence between the full tensor subcategories of $\text{Rep}_f G$ and the quotients G/N , where $N \subset G$ is a closed normal subgroup. Now the inverse limit

$$G = \lim_{\substack{\longleftarrow \\ i \in I}} G_i = \{(g_i \in G_i)_{i \in I} \mid \beta^{i,j}(g_j) = g_i \text{ whenever } i \leq j\}$$

is a compact group with obvious surjective homomorphisms $\gamma_i : G \rightarrow G_i$ for all $i \in I$. Now we define a functor $E : \mathcal{C} \rightarrow \text{Rep}_f G$ as follows: For every $X \in \mathcal{C}$ pick an $i \in I$ such that $X \in \mathcal{C}_i$ and define $F(X) = (E_i(X), \pi_i(X) \circ \gamma_i)$. Clearly this is an object in $\text{Rep}_f G$, and its isomorphism class is independent of the chosen $i \in I$. In this way we obtain a functor from $\mathcal{C} = \lim_{\rightarrow} \mathcal{C}_i$ to $\text{Rep}_f G \cong \lim_{\rightarrow} \text{Rep}_f G_i$ that restricts to equivalences $\mathcal{C}_i \rightarrow \text{Rep}_f G_i$. Thus E is full and faithful. Finally, E is essentially surjective since every finite dimensional representation of $G = \lim_{\leftarrow} G_i$ factors through one of the groups G_i . ■

REMARK 412. In view of Remark 410.3, the preceding proof also shows that every compact group is an inverse limit of compact Lie groups.

B.8 Fiber functors from monoids

Our strategy to proving Theorem 411 will be essentially the one of Deligne [Deligne, 1990], replacing however the algebraic geometry in a symmetric abelian category by fairly elementary commutative categorical algebra. There are already several expositions of this proof [Bichon, 1998; Rosenberg, 2000; Háai, 2002], of which we find [Bichon, 1998] the most useful, see also [Bichon, ND]. However, we will give more details than any of these references, and we provide some further simplifications.

The following result clearly shows the relevance of the notions introduced in Subsection A.6 to our aim of proving Theorem 411:

PROPOSITION 413. *Let \mathcal{C} be a TC^* and $\widehat{\mathcal{C}}$ be a \mathbb{C} -linear strict tensor category containing \mathcal{C} as a full tensor subcategory. Let (Q, m, η) be a monoid in $\widehat{\mathcal{C}}$ satisfying*

- (i) $\dim \text{Hom}_{\widehat{\mathcal{C}}}(\mathbf{1}, Q) = 1$. (I.e., $\text{Hom}_{\widehat{\mathcal{C}}}(\mathbf{1}, Q) = \mathbb{C}\eta$.)
- (ii) For every $X \in \mathcal{C}$, there is $n(X) \in \mathbb{Z}_+$ such that $n(X) \neq 0$ whenever $X \not\cong 0$ and an isomorphism $\alpha_X : (Q \otimes X, m \otimes \text{id}_X) \rightarrow n(X) \cdot (Q, m)$ of Q -modules.

Then the functor $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ defined by

$$E : \mathcal{C} \rightarrow \mathcal{H}, \quad X \mapsto \text{Hom}_{\widehat{\mathcal{C}}}(\mathbf{1}, Q \otimes X),$$

together with

$$(61) \quad E(s)\phi = \text{id}_Q \otimes s \circ \phi, \quad s : X \rightarrow Y, \quad \phi \in \text{Hom}(\mathbf{1}, Q \otimes X)$$

is a faithful (strong) tensor functor and satisfies $\dim_{\mathbb{C}} E(X) = n(X)$.

If $\widehat{\mathcal{C}}$ has a symmetry c w.r.t. which (Q, m, η) is commutative then E is symmetric monoidal w.r.t. the symmetry Σ of $\text{Vect}_{\mathbb{C}}$, i.e. $E(c_{X,Y}) = \Sigma_{E(X),E(Y)}$.

Proof. We have $E(X) = \text{Hom}(\mathbf{1}, Q \otimes X) \cong \text{Hom}(\mathbf{1}, n(X)Q) \cong d(X)\text{Hom}(\mathbf{1}, Q) \cong \mathbb{C}^{n(X)}$, thus $E(X)$ is a vector space of dimension $n(X)$. Since $E(X) \neq 0$ for every non-zero $X \in \mathcal{C}$, the functor E is faithful.

To see that E is monoidal first observe that by (ii) we have $E(\mathbf{1}) = \text{Hom}(\mathbf{1}, Q) = \mathbb{C}\eta$. Thus there is a canonical isomorphism $e : \mathbb{C} = \mathbf{1}_{\text{Vect}_{\mathbb{C}}} \rightarrow E(\mathbf{1}) = \text{Hom}(\mathbf{1}, Q)$ defined by $c \mapsto c\eta$. Next we define morphisms

$$[d_{X,Y}^E : E(X) \otimes E(Y) \rightarrow E(X \otimes Y), \quad \phi \otimes \psi \mapsto m \otimes \text{id}_{X \otimes Y} \circ \text{id}_Q \otimes \phi \otimes \text{id}_Y \circ \psi.$$

By definition (61) of the map $E(s) : E(X) \rightarrow E(Y)$ it is obvious that the family $(d_{X,Y}^E)$ is natural w.r.t. both arguments. The equation

$$d_{X_1 \otimes X_2, X_3}^E \circ d_{X_1, X_2}^E \otimes \text{id}_{E(X_3)} = d_{X_1, X_2 \otimes X_3}^E \circ \text{id}_{E_1} \otimes d_{X_2, X_3}^E \quad \forall X_1, X_2, X_3 \in \mathcal{C}$$

required from a tensor functor is a straightforward consequence of the associativity of m . The verification is left as an exercise.

That $(E, (d_{X,Y}), e)$ satisfies the unit axioms is almost obvious. The first condition follows by

$$d_{X,\mathbf{1}}(\text{id}_{E(X)} \otimes e)\phi = d_{X,\mathbf{1}}(\phi \otimes \eta) = m \otimes \text{id}_X \circ \text{id}_Q \otimes \phi \circ \eta = \phi,$$

and the second is shown analogously.

So far, we have shown that E is a weak tensor functor for which $e : \mathbf{1}_{\mathcal{H}} \rightarrow E(\mathbf{1}_{\mathcal{C}})$ is an isomorphism. In order to conclude that E is a (strong) tensor functor it remains to show that the morphisms $d_{X,Y}^E$ are isomorphisms. Let $X, Y \in \mathcal{C}$. We consider the bilinear map

$$\begin{aligned} \gamma_{X,Y} : \quad & \text{Hom}_Q(Q, Q \otimes X) \boxtimes \text{Hom}_Q(Q, Q \otimes Y) \rightarrow \text{Hom}_Q(Q, Q \otimes X \otimes Y), \\ & s \boxtimes t \mapsto s \otimes \text{id}_Y \circ t. \end{aligned}$$

(We write \boxtimes rather than $\otimes_{\mathbb{C}}$ for the tensor product of $\text{Vect}_{\mathbb{C}}$ in order to avoid confusion with the tensor product in $Q - \text{Mod}$.) By 2., we have Q -module morphisms $s_i : Q \rightarrow Q \otimes X, s'_i : Q \otimes X \rightarrow Q$ for $i = 1, \dots, n(X)$ satisfying $s'_i \circ s_j = \delta_{ij} \text{id}_Q$, and $\sum_i s_i \circ s'_i = \text{id}_{Q \otimes X}$, and similar morphisms $t_i, t'_i, i = 1, \dots, n(Y)$ for X replaced by Y . Then the $\gamma_{ij} = \gamma_{X,Y}(s_i \otimes t_j)$ are linearly independent, since they satisfy $\gamma'_{i'j'} \circ \gamma_{ij} = \delta_{i'i} \delta_{j'j} \text{id}_Q$ with $\gamma'_{i'j'} = t'_{j'} \circ s'_i \otimes \text{id}_Y$. Bijectivity of $\gamma_{X,Y}$ follows now

from the fact that both domain and codomain of $\gamma_{X,Y}$ have dimension $n(X)n(Y)$. Appealing to the isomorphisms $\delta_X : \text{Hom}_Q(Q, Q \otimes X) \mapsto \text{Hom}(\mathbf{1}, Q \otimes X)$ one easily shows

$$d_{X,Y}^E = \delta_{X \otimes Y} \circ \gamma_{X,Y} \circ \delta_X^{-1} \boxtimes \delta_Y^{-1},$$

which implies that $d_{X,Y}^E$ is an isomorphism for every $X, Y \in \mathcal{C}$.


We now assume that $\widehat{\mathcal{C}}$ has a symmetry c and that (Q, m, η) is commutative. In order to show that E is a symmetric tensor functor we must show that

$$E(c_{X,Y}) \circ d_{X,Y}^E = \Sigma_{E(X),E(Y)} \circ d_{Y,X}^E$$

for all $X, Y \in \mathcal{C}$. Let $\phi \in E(X), \psi \in E(Y)$.


By definition of E we have

$$(E(c_{X,Y}) \circ d_{X,Y}^E)(\phi \otimes \psi) = \text{id}_Q \otimes c_{X,Y} \circ m \otimes \text{id}_{X \otimes Y} \circ \text{id}_Q \otimes \phi \otimes \text{id}_Y \circ \psi$$



On the other hand,

$$(d_{Y,X}^E \circ c_{E(X),E(Y)})(\phi \otimes \psi) = (d_{Y,X}^E \circ \Sigma_{E(X),E(Y)})(\phi \otimes \psi) = d_{Y,X}^E(\psi \otimes \phi)$$



If m is commutative, i.e. $m = m \circ c_{Q,Q}$, these two expressions coincide, and we are done. ■

REMARK 414. 1. The property (ii) in the proposition is called the ‘absorbing property’.

2. The conditions in Proposition 413 are in fact necessary for the existence of a fiber functor! Assume that a tensor $*$ -category \mathcal{C} admits a $*$ -preserving fiber functor $E : \mathcal{C} \rightarrow \mathcal{H}$. By [Müger *et al.*, 2004], which reviews and extends work of Woronowicz, Yamagami and others, there is a discrete algebraic quantum group (A, Δ) such that $\mathcal{C} \simeq \text{Rep}_f(A, \Delta)$. In [Müger and Tuset, 2006] it is shown that taking $\widehat{\mathcal{C}} \simeq \text{Rep}(A, \Delta)$ (i.e. representations of any dimension) and $Q = \pi_l$, there is a monoid (Q, m, η) satisfying the conditions of Proposition 413. Namely, one can take $Q = \pi_l$, the left regular representation. In [Müger and Tuset, 2006] it is shown that (i) $\dim \text{Hom}(\pi_0, \pi_l) = 1$, i.e. there exists a non-zero morphism $\eta : \pi_0 \rightarrow \pi_l$, unique up to normalization; (ii) π_l has the required absorbing property; (iii) there exists a morphism $m : \pi_l \otimes \pi_l \rightarrow \pi_l$ such that $(Q = \pi_l, m, \eta)$ is a monoid.

3. In the previous situation, the left regular representation π_l lives in $\text{Rep}_f(A, \Delta)$ iff A is finite dimensional. This already suggests that the category \mathcal{C} in general is too small to contain a monoid of the desired properties. In fact, assume we can take $\widehat{\mathcal{C}} = \mathcal{C}$. Then for every irreducible $X \in \mathcal{C}$ we have $\dim \text{Hom}(X, Q) = \dim \text{Hom}(\mathbf{1}, Q \otimes \overline{X}) = n(\overline{X}) > 0$. Thus Q contains all irreducible objects as direct summands. Since every object in \mathcal{C} is a finite direct sum of simple objects, $\widehat{\mathcal{C}} = \mathcal{C}$ is possible only if \mathcal{C} has only finitely many isomorphism classes of simple objects. In fact, even in this case, our construction of (Q, m, η) will require the use of a bigger category $\widehat{\mathcal{C}}$. It is here that the category $\text{Ind } \mathcal{C}$ of Subsection A.7 comes into play.

Since we have already reduced the problem of constructing a fiber functor to the case of finitely generated tensor categories, we want a version of the preceding result adapted to that situation:

COROLLARY 415. *Let \mathcal{C} be a TC^* with monoidal generator $Z \in \mathcal{C}$ and let $\widehat{\mathcal{C}}$ be a \mathbb{C} -linear strict tensor category containing \mathcal{C} as a full tensor subcategory. If (Q, m, η) is a monoid in $\widehat{\mathcal{C}}$ satisfying*

(i) $\dim \text{Hom}_{\widehat{\mathcal{C}}}(\mathbf{1}, Q) = 1$.

(ii) *There is $d \in \mathbb{N}$ and an isomorphism $\alpha_Z : (Q \otimes Z, m \otimes \text{id}_Z) \rightarrow d \cdot (Q, m)$ of Q -modules.*

Then the hypothesis (ii) in Proposition 413 follows. Thus $E : X \mapsto \text{Hom}_{\widehat{\mathcal{C}}}(\mathbf{1}, Q \otimes X)$ is a fiber functor.

Proof. If $X \in \mathcal{C}$, there exists $n \in \mathbb{N}$ such that $X \prec Z^{\otimes n}$. Concretely, there are morphisms $u : X \rightarrow Z^{\otimes n}$ and $v : Z^{\otimes n} \rightarrow X$ such that $v \circ u = \text{id}_X$. Then the morphisms $\tilde{u} = \text{id}_Q \otimes u : Q \otimes X \rightarrow Q \otimes Z^{\otimes n}$ and $\tilde{v} = \text{id}_Q \otimes v : Q \otimes Z^{\otimes n} \rightarrow Q \otimes X$ are morphisms of Q -modules. Thus the Q -module $(Q \otimes X, m \otimes \text{id}_X)$ is a direct summand of $(Q \otimes Z^{\otimes n}, m \otimes \text{id}_{Z^{\otimes n}})$. By assumption, the latter is isomorphic to a direct sum of d^n copies of (Q, m) . By Lemma 358 and assumption

(i), $\text{End}_{\mathcal{Q}}((Q, m)) \cong \mathbb{C}$, thus $(Q, m) \in \mathcal{Q} - \text{Mod}$ is irreducible. Thus the direct summand $(Q \otimes X, m \otimes \text{id}_X)$ of $d^n \cdot (Q, m)$ is a direct sum of r copies of (Q, m) with $r \leq d^m$ and $r \neq 0$ whenever $X \neq 0$. Thus hypothesis (ii) in Proposition 413 holds. ■

In view of Corollary 415, proving Theorem 411 amounts to finding a symmetric tensor category $\widehat{\mathcal{C}}$ containing \mathcal{C} as a full subcategory and a commutative monoid (Q, m, η) in $\widehat{\mathcal{C}}$ such that $\dim \text{Hom}(\mathbf{1}, Q) = 1$ and $Q \otimes Z \cong d \otimes Q$ as Q -modules for a suitable monoidal generator Z of \mathcal{C} . This will be achieved in Subsection B.11, based on thorough analysis of the permutation symmetry of the category \mathcal{C} .

B.9 Symmetric group action, determinants and integrality of dimensions

We now turn to a discussion of certain representations of the symmetric groups $P_n, n \in \mathbb{N}$, present in tensor $*$ -categories with a unitary symmetry. It is well known that the symmetric group P_n on n labels has the presentation

$$P_n = (\sigma_1, \dots, \sigma_{n-1} \mid |i - j| \geq 2 \Rightarrow \sigma_i \sigma_j = \sigma_j \sigma_i, \\ \sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1} \ \forall i \in \{1, \dots, n - 1\}, \ \sigma_i^2 = 1 \ \forall i).$$

Since \mathcal{C} is strict we may define the tensor powers $X^{\otimes n}, n \in \mathbb{N}$, in the obvious way for any $X \in \mathcal{C}$. We posit $X^{\otimes 0} = \mathbf{1}$ for every $X \in \mathcal{C}$.

LEMMA 416. *Let \mathcal{C} be an STC*. Let $X \in \mathcal{C}$ and $n \in \mathbb{N}$. Then*

$$\Pi_n^X : \sigma_i \mapsto \text{id}_{X^{\otimes i-1}} \otimes c_{X,X} \otimes \text{id}_{X^{\otimes n-i-1}}$$

uniquely determines a homomorphism Π_n^X from the group P_n into the unitary group of $\text{End } X^{\otimes n}$.

Proof. It is clear that $\Pi_n^X(\sigma_i)$ and $\Pi_n^X(\sigma_j)$ commute if $|i - j| \geq 2$. That $\Pi_n^X(\sigma_i)^2 = \text{id}_{X^{\otimes n}}$ is equally obvious. Finally,

$$\Pi_n^X(\sigma_i) \circ \Pi_n^X(\sigma_{i+1}) \circ \Pi_n^X(\sigma_i) = \Pi_n^X(\sigma_{i+1}) \circ \Pi_n^X(\sigma_i) \circ \Pi_n^X(\sigma_{i+1})$$

follows from the Yang-Baxter equation satisfied by the symmetry c . ■

REMARK 417. Dropping the relations $\sigma_i^2 = 1$ the same formulae as above define homomorphisms of the Artin braid groups B_n into $\text{End } X^{\otimes n}$. However, none of the following considerations has known analogues in the braided case.

Recall that there is a homomorphism $\text{sgn} : P_n \rightarrow \{1, -1\}$, the signature map.

LEMMA 418. *Let \mathcal{C} be an STC*. For any $X \in \mathcal{C}$ we define orthogonal projections in $\text{End } X^{\otimes 0} = \text{End } \mathbf{1}$ by $S_0^X = A_0^X = \text{id}_{\mathbf{1}}$. For any $n \in \mathbb{N}$, the morphisms*

$$S_n^X = \frac{1}{n!} \sum_{\sigma \in P_n} \Pi_n^X(\sigma), \\ A_n^X = \frac{1}{n!} \sum_{\sigma \in P_n} \text{sgn}(\sigma) \Pi_n^X(\sigma)$$

satisfy

$$\begin{aligned} \Pi_n^X(\sigma) \circ S_n^X &= S_n^X \circ \Pi_n^X(\sigma) = S_n^X, \\ \Pi_n^X(\sigma) \circ A_n^X &= A_n^X \circ \Pi_n^X(\sigma) = \text{sgn}(\sigma) A_n^X \end{aligned}$$

for all $\sigma \in P_n$ and are thus orthogonal projections in the $*$ -algebra $\text{End } X^{\otimes n}$.

Proof. Straightforward computations. ■

DEFINITION 419. The subobjects (defined up to isomorphism) of $X^{\otimes n}$ corresponding to the idempotents S_n^X and A_n^X are denoted by $S_n(X)$ and $A_n(X)$, respectively.

The following was proven both in [Doplicher and Roberts, 1989] and [Deligne, 1990]:

PROPOSITION 420. *Let \mathcal{C} be an even STC*. For any $X \in \mathcal{C}$ we have*

$$(62) \quad \text{Tr}_{X^{\otimes n}} A_n^X = \frac{d(X)(d(X) - 1)(d(X) - 2) \cdots (d(X) - n + 1)}{n!} \quad \forall n \in \mathbb{N}.$$

Proof. (*Sketch*) Making crucial use of the fact that \mathcal{C} is even, i.e. $\Theta(X) = \text{id}_X$ for all $X \in \mathcal{C}$, one can prove

$$\text{Tr}_{X^{\otimes n}} \Pi_n^X(\sigma) = d(X)^{\#\sigma} \quad \forall X \in \mathcal{C}, \sigma \in P_n,$$

where $\#\sigma$ is the number of cycles into which the permutation σ decomposes. (The reader familiar with tangle diagrams will find this formula almost obvious: Triviality of the twist $\Theta(X)$ implies invariance under the first Reidemeister move. Thus the closure of the permutation σ is equivalent to $\#\sigma$ circles, each of which contributes a factor $d(X)$.) Now the result follows at once from the definition of A_n^X and the formula

$$\sum_{\sigma \in P_n} \text{sgn}(\sigma) z^{\#\sigma} = z(z - 1)(z - 2) \cdots (z - n + 1),$$

which holds for all $n \in \mathbb{N}$ and $z \in \mathbb{C}$, as one can prove by induction over n . ■

COROLLARY 421. *In an STC* we have $d(X) \in \mathbb{N}$ for every non-zero $X \in \mathcal{C}$.*

Proof. Assume first that \mathcal{C} is even, and let $X \in \mathcal{C}$. Since \mathcal{C} has subobjects there exist an object $A_n(X) \in \mathcal{C}$ and a morphism $s : A_n(X) \rightarrow X^{\otimes n}$ such that $s^* \circ s = \text{id}_{A_n(X)}$ and $s \circ s^* = A_n^X$. Then by part 1 and 2 in Proposition 339, we get

$$\text{Tr}_{X^{\otimes n}} A_n^X = \text{Tr}_{X^{\otimes n}}(s \circ s^*) = \text{Tr}_{A_n(X)}(s^* \circ s) = \text{Tr}_{A_n(X)} \text{id}_{A_n(X)} = d(A_n(X)).$$

Since the dimension of any object in a $*$ -category is non-negative we thus conclude that $\text{Tr}_{X^{\otimes n}} A_n^X \geq 0$ for all $n \in \mathbb{N}$. From the right-hand side in the formula (62)

for $Tr_{X^{\otimes n}} A_n^X$ we see that $Tr_{X^{\otimes n}} A_n^X$ will become negative for some $n \in \mathbb{N}$ unless $d(X) \in \mathbb{N}$.

If \mathcal{C} is odd, the above argument gives integrality of the dimensions in the bosonized category $\tilde{\mathcal{C}}$. Since the categorical dimension is independent of the braiding, we have $d_{\mathcal{C}}(X) = d_{\tilde{\mathcal{C}}}(X)$ and are done. ■

Let \mathcal{C} be an STC^* and $X \in \mathcal{C}$ non-zero and set $d = d(X) \in \mathbb{N}$. Consider the subobject $A_d(X)$ of $X^{\otimes d}$, introduced in the proof of Corollary 421, which corresponds to the orthogonal projection $A_d^X \in \text{End } X^{\otimes d}$ defined in Lemma 418. Then

$$d(A_d(X)) = Tr_{X^{\otimes d}} A_d^X = \frac{d!}{d!} = 1,$$

we see that $A_d(X)$ is an irreducible and invertible object of \mathcal{C} (with inverse $\overline{A_d(X)}$).

DEFINITION 422. The isomorphism class of $A^{d(X)}(X)$ is called the *determinant* $\det X$ of X .

LEMMA 423. Let \mathcal{C} be an STC^* and $X, Y \in \mathcal{C}$. Then

- (i) $\det \overline{X} \cong \overline{\det(X)}$.
- (ii) $\det(X \oplus Y) \cong \det X \otimes \det Y$.
- (iii) $\det(X \oplus \overline{X}) \cong \mathbf{1}$.

Proof. (i) Let $(\overline{X}, r, \bar{r})$ be a standard left inverse of X . By inductive use of Lemma 338 one obtains standard left inverses $(\overline{X}^{\otimes n}, r_n, \bar{r}_n)$ of $X^{\otimes n}$ for any $n \in \mathbb{N}$. If now $\sigma = \sigma_{i_1} \cdots \sigma_{i_r} \in P_n$, one can verify that

$$\Pi_n^{\overline{X}}(\sigma') = r_n^* \otimes \text{id}_{\overline{X}^{\otimes n}} \circ \text{id}_{\overline{X}^{\otimes n}} \otimes \Pi_n^X(\sigma) \otimes \text{id}_{\overline{X}^{\otimes n}} \circ \text{id}_{\overline{X}^{\otimes n}} \otimes \bar{r}_n,$$

where $\sigma' = \sigma_{n-i_r}^{-1} \cdots \sigma_{n-i_1}^{-1}$. In particular, $\text{sgn } \sigma' = \text{sgn } \sigma$, implying

$$A_n^{\overline{X}} = r_n^* \otimes \text{id}_{\overline{X}^{\otimes n}} \circ \text{id}_{\overline{X}^{\otimes n}} \otimes A_n^X \otimes \text{id}_{\overline{X}^{\otimes n}} \circ \text{id}_{\overline{X}^{\otimes n}} \otimes \bar{r}_n,$$

for any $n \in \mathbb{N}$. Now the claim follows from Lemma 337.

(ii) For any $X \in \mathcal{C}$ we abbreviate $d_X = d(X)$ and $A^X = A_{d_X}^X \in \text{End } X^{\otimes d_X}$. Let $u : X \rightarrow Z, v : Y \rightarrow Z$ be isometries implementing $Z \cong X \oplus Y$. Then $X^{\otimes d_X}$ is a subobject of $Z^{\otimes d_X}$, and similarly for $Y^{\otimes d_Y}$. By definition, $\det Z$ is the subobject of $Z^{\otimes d_Z}$ corresponding to the projector $A^Z \in \text{End } Z^{\otimes d_Z}$. On the other hand, $\det X \otimes \det Y$ is the subobject of $X^{\otimes d_X} \otimes Y^{\otimes d_Y}$ corresponding to the projector $A^X \otimes A^Y$, and therefore it is isomorphic to the subobject of $Z^{\otimes d_Z}$ corresponding to the projector

$$u \otimes \cdots \otimes u \otimes v \otimes \cdots \otimes v \circ A^X \otimes A^Y \circ u^* \otimes \cdots \otimes u^* \otimes v^* \otimes \cdots \otimes v^* \in \text{End } Z^{\otimes d_Z},$$

where there are d_X factors u and u^* and d_Y factors v and v^* . This equals

$$\frac{1}{d_X!d_Y!} \sum_{\substack{\sigma \in P_{d_X} \\ \sigma' \in P_{d_Y}}} \text{sgn}(\sigma)\text{sgn}(\sigma') u \otimes \cdots \otimes u \otimes v \otimes \cdots \otimes v \\ \circ \Pi_{d_X}^X(\sigma) \otimes \Pi_{d_Y}^Y(\sigma') \circ u^* \otimes \cdots \otimes u^* \otimes v^* \otimes \cdots \otimes v^*.$$

By naturality of the braiding, this equals

$$\frac{1}{d_X!d_Y!} \sum_{\substack{\sigma \in P_{d_X} \\ \sigma' \in P_{d_Y}}} \text{sgn}(\sigma)\text{sgn}(\sigma') \Pi_{d_X}^Z(\sigma) \otimes \Pi_{d_Y}^Z(\sigma') \circ p_X \otimes \cdots \otimes p_X \otimes p_Y \otimes \cdots \otimes p_Y,$$

where $p_X = u \circ u^*, p_Y = v \circ v^*$. With the juxtaposition $\sigma \times \sigma' \in P_{d_X+d_Y} = P_{d_Z}$ of σ and σ' this becomes

$$(63) \quad \frac{1}{d_X!d_Y!} \sum_{\substack{\sigma \in P_{d_X} \\ \sigma' \in P_{d_Y}}} \text{sgn}(\sigma \times \sigma') \Pi_{d_Z}^Z(\sigma \times \sigma') \circ p_X \otimes \cdots \otimes p_X \otimes p_Y \otimes \cdots \otimes p_Y.$$

On the other hand,

$$\begin{aligned} A^Z &= \frac{1}{d_Z!} \sum_{\sigma \in P_{d_Z}} \text{sgn}(\sigma) \Pi_{d_Z}^Z(\sigma) \\ &= \left(\sum_{\sigma \in P_{d_Z}} \text{sgn}(\sigma) \Pi_{d_Z}^Z(\sigma) \right) \circ (p_X + p_Y) \otimes \cdots \otimes (p_X + p_Y). \end{aligned}$$

Of the 2^{d_Z} terms into which this can be decomposed, only those with d_X factors p_X and d_Y factors p_Y are nonzero since $A_n^X = 0$ for $n > d_X$ and $A_n^Y = 0$ for $n > d_Y$. We are thus left with a sum of $d_Z!/d_X!d_Y!$ terms, and working out the signs we see that they all equal to $d_X!d_Y!/d_Z!$ times (63), thus the sum equals (63). This proves the isomorphism $\det Z \cong \det X \otimes \det Y$.

Finally, (iii) follows from

$$\det(X \oplus \overline{X}) \cong \det X \otimes \det \overline{X} \cong \det X \otimes \overline{\det X} \cong \det X \otimes (\det X)^{-1} \cong \mathbf{1},$$

where we have used (i) and (ii) of this lemma, $d(\det X) = 1$ and (iii) of Lemma 341. ■

For later use we state a computational result:

LEMMA 424. *Let X satisfy $\det X \cong \mathbf{1}$ and write $d = d(X)$. If $s : \mathbf{1} \rightarrow X^{\otimes d}$ is an isometry for which $s \circ s^* = A_d^X$ then*

$$(64) \quad s^* \otimes \text{id}_X \circ \text{id}_X \otimes s = (-1)^{d-1} d^{-1} \text{id}_X.$$

Proof. We abbreviate $x = s^* \otimes \text{id}_X \circ \text{id}_X \otimes s$ and observe that by non-degeneracy of the trace it is sufficient to show that $\text{Tr}_X(ax) = (-1)^{d-1}d^{-1}\text{Tr}_X(a)$ for all $a \in \text{End}X$. In order to show this, let (\bar{X}, r, \bar{r}) be a standard solution of the conjugate equations and compute

$$\begin{aligned}
 \text{Tr}_X(ax) &= \bar{X} \text{ (diagram with } s^*, a, s, X, X^{d-1} \text{)} = (-1)^{d-1} \bar{X} \text{ (diagram with } s^*, a, s, X, X^{d-1} \text{)} \\
 &= (-1)^{d-1} \bar{X} \text{ (diagram with } s^*, a, s, X, X^{d-1} \text{)} = (-1)^{d-1} \bar{X} \text{ (diagram with } s^*, a, s, X, X^{d-1} \text{)}
 \end{aligned}$$

We have in turn used the total antisymmetry of s (Lemma 418), the naturality properties of the braiding and the triviality of the twist Θ_X . Now,

$$\begin{aligned}
 s^* \circ a \otimes \text{id}_{X^{\otimes d-1}} \circ s &= \text{Tr}_1(s^* \circ a \otimes \text{id}_{X^{\otimes d-1}} \circ s) \\
 &= \text{Tr}_{X^{\otimes d}}(a \otimes \text{id}_{X^{\otimes d-1}} \circ s \circ s^*) = \text{Tr}_{X^{\otimes d}}(a \otimes \text{id}_{X^{\otimes d-1}} \circ A_d^X).
 \end{aligned}$$

In order to complete the proof we need to show that this equals $d^{-1}\text{Tr}_X a$, which is done by suitably modifying the proof of Proposition 420. By the same argument as given there, it suffices to prove $\text{Tr}_{X^{\otimes d}}(a \otimes \text{id}_{X^{\otimes d-1}} \circ \Pi_d^X(\sigma)) = d^{\#\sigma-1}\text{Tr}_X a$. Again, the permutation σ decomposes into a set of cyclic permutations, of which now precisely one involves the index 1. It is therefore sufficient to prove $\text{Tr}_{X^{\otimes n}}(a \otimes \text{id}_{X^{\otimes n-1}} \circ \Pi_n^X(\sigma)) = \text{Tr}_X a$ for every cyclic permutation σ of all n indices. Inserting a at the appropriate place, the calculation essentially proceeds as before. The only difference is that instead of $\text{Tr}_X \text{id}_X = d(X)$ one is left with $\text{Tr}_X a$, giving rise to the desired result. ■

REMARK 425. Objects with determinant **1** were called special in [Doplicher and Roberts, 1989], where also all results of this subsection can be found.

This concludes our discussion of antisymmetrization and determinants, and we turn to symmetrization and the symmetric algebra. It is here that we need the Ind-category that was introduced in Subsection A.7.

B.10 The symmetric algebra

In “ordinary” algebra one defines the symmetric algebra $S(V)$ over a vector space V . Unless $V = \{0\}$, this is an infinite direct sum of non-trivial vector spaces. We will need a generalization of this construction to symmetric tensor categories other than Vect. While infinite direct sums of objects make sense in the setting of C^* -tensor categories (Definition 345), a more convenient setting for the following considerations is given by the theory of abelian categories.

LEMMA 426. *Let \mathcal{C} be an STC^* and $X \in \mathcal{C}$. For every $n \in \mathbb{N}$ choose an object $S_n(X)$ and an isometry $u_n : S_n(X) \rightarrow X^{\otimes n}$ such that $u_n \circ u_n^* = S_n^X$. Also, let $u_0 = \text{id}_{\mathbf{1}}$, interpreted as a morphism from $S_0(X) = \mathbf{1}$ to $X^0 = \mathbf{1}$. The morphisms $m_{i,j} : S_i(X) \otimes S_j(X) \rightarrow S_{i+j}(X)$ defined by*

$$m_{i,j} : S_i(X) \otimes S_j(X) \xrightarrow{u_i \otimes u_j} X^{\otimes i} \otimes X^{\otimes j} \equiv X^{\otimes(i+j)} \xrightarrow{u_{i+j}^*} S_{i+j}(X)$$

satisfy

$$m_{i+j,k} \circ m_{i,j} \otimes \text{id}_{S_k(X)} = m_{i,j+k} \circ \text{id}_{S_i(X)} \otimes m_{j,k}$$

for all $i, j, k \in \mathbb{Z}_+$. Furthermore,

$$m_{i,j} = m_{j,i} \circ c_{S_i(X), S_j(X)} \quad \forall i, j$$

and $m_{i,0} = m_{0,i} = \text{id}_{S_i(X)}$.

Proof. As a consequence of $S_n^X \circ \Pi_n^X(\sigma) = S_n^X(\sigma)$ for all $\sigma \in P_n$, cf. Lemma 418, we have

$$S_{i+j+k}^X \circ S_{i+j}^X \otimes \text{id}_{X^{\otimes k}} \circ S_i^X \otimes S_j^X \otimes \text{id}_{X^{\otimes k}} = S_{i+j+k}^X \circ S_{i+j}^X \otimes \text{id}_{X^{\otimes k}} = S_{i+j+k}^X,$$

$$S_{i+j+k}^X \circ \text{id}_{X^{\otimes k}} \otimes S_{j+k}^X \circ \text{id}_{X^{\otimes k}} \otimes S_j^X \otimes S_k^X = S_{i+j+k}^X \circ \text{id}_{X^{\otimes i}} \otimes S_{j+k}^X = S_{i+j+k}^X.$$

Multiplying all this with u_{i+j+k}^* on the left and with $u_i \otimes u_j \otimes u_k$ on the right and using $u_i^* \circ S_i^X = u_n^*$ and $S_i^X \circ u_i = u_i$ this implies

$$\begin{aligned} u_{i+j+k}^* \circ S_{i+j}^X \otimes \text{id}_{X^{\otimes k}} \circ u_i \otimes u_j \otimes u_k &= u_{i+j+k}^* \circ u_i \otimes u_j \otimes u_k \\ &= u_{i+j+k}^* \circ \text{id}_{X^{\otimes k}} \otimes S_{j+k}^X \circ u_i \otimes u_j \otimes u_k. \end{aligned}$$

Using again that $S_{i+j}^X = u_{i+j} \circ u_{i+j}^*$, we have the first identity we wanted to prove. Furthermore,

$$\begin{aligned} m_{j,i} \circ c_{S_i(X), S_j(X)} &= u_{i+j}^* \circ u_j \otimes u_i \circ c_{S_i(X), S_j(X)} = u_{i+j}^* \circ c_{X^{\otimes i}, X^{\otimes j}} \circ u_i \otimes u_j \\ &= u_{i+j}^* \circ \Pi_{i+j}^X(\sigma) \circ u_i \otimes u_j = u_{i+j}^* \circ S_{i+j}^X \circ \Pi_{i+j}^X(\sigma) \circ u_i \otimes u_j \\ &= u_{i+j}^* \circ S_{i+j}^X \circ u_i \otimes u_j = u_{i+j}^* \circ u_i \otimes u_j = m_{i,j}, \end{aligned}$$

where $\sigma \in P_{i+j}$ is the permutation exchanging the first i with the remaining j strands. The last claim is obvious in view of $S_0(X) = \mathbf{1}$. ■

In view of Lemma 353, \mathcal{C} (with a zero object thrown in) is an abelian category, thus there exists an abelian \mathbb{C} -linear strict symmetric tensor category $\text{Ind } \mathcal{C}$ containing \mathcal{C} as a full subcategory and complete w.r.t. filtered inductive limits. Therefore, for any object X in the STC^* \mathcal{C} , there exists an object

$$S(X) = \varinjlim_{n \rightarrow \infty} \bigoplus_{i=0}^n S_n(X)$$

together with monomorphisms $v_n : S_n(X) \rightarrow S(X)$.

PROPOSITION 427. *Let \mathcal{C} be an STC^* and $X \in \mathcal{C}$. Then there exists a morphism $m_{S(X)} : S(X) \otimes S(X) \rightarrow S(X)$ such that*

$$m_{S(X)} \circ v_i \otimes v_j = v_{i+j} \circ m_{i,j} : S_i(X) \otimes S_j(X) \rightarrow S(X)$$

and $(S(X), m_{S(X)}, \eta_{S(X)} \equiv v_0)$ is a commutative monoid in $\text{Ind } \mathcal{C}$.

Proof. This amounts to using

$$\text{Hom}_{\text{Ind } \mathcal{C}}(S(X) \otimes S(X), S(X)) = \varinjlim_m \varinjlim_n \text{Hom}_{\mathcal{C}} \left(\bigoplus_{i,j=0}^m S_i(X) \otimes S_j(X), \bigoplus_{k=0}^n S_k(X) \right)$$

to assemble the morphisms $m_{i,j} : S_i(X) \otimes S_j(X) \rightarrow S_{i+j}(X)$ into one big morphism $S(X) \otimes S(X) \rightarrow S(X)$. We omit the tedious but straightforward details. Associativity ($m_{S(X)} \circ m_{S(X)} \otimes \text{id}_{S(X)} = m_{S(X)} \circ \text{id}_{S(X)} \otimes m_{S(X)}$) and commutativity ($m_{S(X)} = m_{S(X)} \circ c_{S(X), S(X)}$) then follow from the respective properties of the $m_{i,j}$ established in Lemma 426. The unit property $m_{S(X)} \circ \text{id}_{S(X)} \otimes v_0 = \text{id}_{S(X)} \otimes v_0 = \text{id}_{S(X)}$ follows from $m_{i,0} = m_{0,i} = \text{id}_{S_i(X)}$. ■

We now study the interaction between the operations of symmetrization and antisymmetrization, i.e. between determinants and symmetric algebras, that lies at the core of the embedding theorem. We begin by noting that given two commutative monoids (Q_i, m_i, η_i) , $i = 1, 2$ in a strict symmetric tensor category, the triple $(Q_1 \otimes Q_2, m_{Q_1 \otimes Q_2}, \eta_{Q_1 \otimes Q_2})$, where $\eta_{Q_1 \otimes Q_2} = \eta_1 \otimes \eta_2$ and

$$m_{Q_1 \otimes Q_2} = m_1 \otimes m_2 \circ \text{id}_{Q_1} \otimes c_{Q_2, Q_1} \otimes \text{id}_{Q_2},$$

defines a commutative monoid, the direct product $(Q_1, m_1, \eta_1) \times (Q_2, m_2, \eta_2)$. The direct product \times is strictly associative, thus multiple direct products are unambiguously defined by induction.

LEMMA 428. *Let \mathcal{C} be a STC and assume $Z \in \mathcal{C}$ satisfies $\det Z \cong \mathbf{1}$. We write $d = d(Z)$ and pick $s : \mathbf{1} \rightarrow Z^{\otimes d}$, $s' : Z^{\otimes d} \rightarrow \mathbf{1}$ such that $s' \circ s = \text{id}_{\mathbf{1}}$ and $s \circ s' = A_d^Z$. Let $S(Z)$ be the symmetric tensor algebra over Z with the canonical embeddings $v_0 : \mathbf{1} \rightarrow S(Z), v_1 : Z \rightarrow S(Z)$. Consider the commutative monoid structure on $Q = S(Z)^{\otimes d}$ given by*

$$(Q, m_Q, \eta_Q) = (S(Z), m_{S(Z)}, \eta_{S(Z)})^{\times d}.$$

Define morphisms $f : \mathbf{1} \rightarrow Q$ and $u_i : Z \rightarrow Q$, $t_i : Z^{\otimes(d-1)} \rightarrow Q$, $i = 1, \dots, d$ by

$$f = \underbrace{v_1 \otimes \dots \otimes v_1}_{d \text{ factors}} \circ s,$$

$$u_i = \underbrace{v_0 \otimes \dots \otimes v_0}_{i-1 \text{ factors}} \otimes v_1 \otimes \underbrace{v_0 \otimes \dots \otimes v_0}_{d-i \text{ factors}},$$

$$t_i = (-1)^{d-i} \underbrace{v_1 \otimes \dots \otimes v_1}_{i-1 \text{ factors}} \otimes v_0 \otimes \underbrace{v_1 \otimes \dots \otimes v_1}_{d-i \text{ factors}}.$$

Then s, f, u_i, t_j satisfy

$$(65) \quad m_Q \circ t_j \otimes u_i \circ s = \delta_{ij} f \quad \forall i, j \in \{1, \dots, d\}.$$

Proof. First note that $s : \mathbf{1} \rightarrow Z^{\otimes d}$ as required exists since $\det Z \cong \mathbf{1}$ and that f is a composition of monics, thus non-zero. We compute

$$\begin{aligned} m_Q \circ t_i \otimes u_i \circ s &= (-1)^{d-i} \text{id}_{S(Z)^{\otimes(i-1)}} \otimes c_{S(Z)^{\otimes(d-i)}, S(Z)} \circ v_1 \otimes v_1 \otimes \dots \otimes v_1 \circ s \\ &= (-1)^{d-i} v_1 \otimes v_1 \otimes \dots \otimes v_1 \circ \text{id}_{Z^{\otimes(i-1)}} \otimes c_{Z^{\otimes(d-i)}, Z} \circ s \\ &= v_1 \otimes v_1 \otimes \dots \otimes v_1 \circ s \\ &= f. \end{aligned}$$

In the first equality we used the definition of (Q, m_Q, η_Q) as d -fold direct product of $(S(Z), m_{S(Z)}, \eta_{S(Z)})$ and the fact that $v_0 = \eta_{S(Z)}$ is the unit, naturality of the braiding in the second and Lemma 418 in the third. To see that $m_Q \circ t_j \otimes u_i \circ s = 0$ if $i \neq j$ consider $j = d - 1, i = d$. Then $m_Q \circ t_j \otimes u_i \circ s : \mathbf{1} \rightarrow S(Z)^{\otimes d} \equiv Q$ is the composite

$$\mathbf{1} \xrightarrow{s} Z^{\otimes d} \xrightarrow{\underbrace{v_1 \otimes \dots \otimes v_1}_{d-2 \text{ factors}} \otimes v_0 \otimes v_1 \otimes v_1} S(Z)^{\otimes(d+1)} \xrightarrow{\text{id}_{S(Z)^{\otimes(d-1)}} \otimes m_{S(Z)}} S(Z)^{\otimes d}.$$

Now,

$$\begin{aligned}
 & \text{id}_{S(Z)^{\otimes(d-1)}} \otimes m_{S(Z)} \circ v_1 \otimes \cdots \otimes v_1 \otimes v_0 \otimes v_1 \otimes v_1 \circ s \\
 &= \text{id}_{S(Z)^{\otimes(d-1)}} \otimes (m_{S(Z)} \circ c_{S(Z),S(Z)}) \circ v_1 \otimes \cdots \otimes v_1 \otimes v_0 \otimes v_1 \otimes v_1 \circ s \\
 &= \text{id}_{S(Z)^{\otimes(d-1)}} \otimes m_{S(Z)} \circ \text{id}_{S(Z)^{\otimes(d-1)}} \otimes c_{S(Z),S(Z)} \\
 &\quad \circ v_1 \otimes \cdots \otimes v_1 \otimes v_0 \otimes v_1 \otimes v_1 \circ s \\
 &= \text{id}_{S(Z)^{\otimes(d-1)}} \otimes m_{S(Z)} \circ v_1 \otimes \cdots \otimes v_1 \otimes v_0 \otimes v_1 \otimes v_1 \circ \text{id}_{Z^{\otimes(d-2)}} \otimes c_{Z,Z} \circ s \\
 &= - \text{id}_{S(Z)^{\otimes(d-1)}} \otimes m_{S(Z)} \circ v_1 \otimes \cdots \otimes v_1 \otimes v_0 \otimes v_1 \otimes v_1 \circ s,
 \end{aligned}$$

where we used the commutativity of $m_{S(Z)}$ in the first step and the total antisymmetry of s in the last. Thus $m_Q \circ u_d \otimes t_{d-1} \circ s = -m_Q \circ u_d \otimes t_{d-1} \circ s = 0$. For general $i \neq j$ the argument is exactly the same, but becomes rather tedious to write up in detail. ■

REMARK 429. Lemma 428 and Proposition 430 below, both taken from [Bichon, 1998], are the crucial ingredients in our approach to the reconstruction theorem.

B.11 Construction of an absorbing commutative monoid

Throughout this subsection, let \mathcal{C} be an even STC^* with monoidal generator Z . Consider the commutative monoid $(Q, m, \eta) = (S(Z), m_{S(Z)}, \eta_{S(Z)})^{\times d(Z)}$ in $\text{Ind } \mathcal{C}$ and the morphisms s, s', f, u_i, t_j as defined in Lemma 428. Then $m_0 \in \text{End } Q$ defined by

$$m_0 = m_Q \circ \text{id}_Q \otimes (f - \eta_Q) = m_Q \circ \text{id}_Q \otimes f - \text{id}_Q$$

is a Q -module map, thus $m_0 \in \text{End}_Q((Q, m_Q))$. Then its image $j = \text{im } m_0 : (J, \mu_J) \rightarrow (Q, m_Q)$ (in the abelian category $Q\text{-Mod}$) defines an ideal $j : (J, \mu_J) \rightarrow (Q, m)$ in (Q, m, η) . This ideal is proper iff j is not an isomorphism iff m_0 is not an isomorphism. Postponing this issue for a minute, we have:

PROPOSITION 430. *Let \mathcal{C} be an even symmetric STC^* and let $Z \in \mathcal{C}$ be such that $\det Z \cong \mathbf{1}$. Let (Q, m, η) and s, s', f, u_i, t_j be as defined in Lemma 428 and m_0 as above. Let $j' : (J', \mu') \rightarrow (Q, m)$ be any proper ideal in (Q, m, η) containing the ideal $j : (J, \mu) \rightarrow (Q, m)$, where $j = \text{im } m_0$. Let (B, m_B, η_B) be the quotient monoid. Then there is an isomorphism*

$$(B \otimes Z, m \otimes \text{id}_Z) \cong d(Z) \cdot (B, m_B)$$

of B -modules.

Proof. Since the ideal is proper, the quotient (B, m_B, η_B) is nontrivial and we have an epi $p : Q \rightarrow B$ satisfying

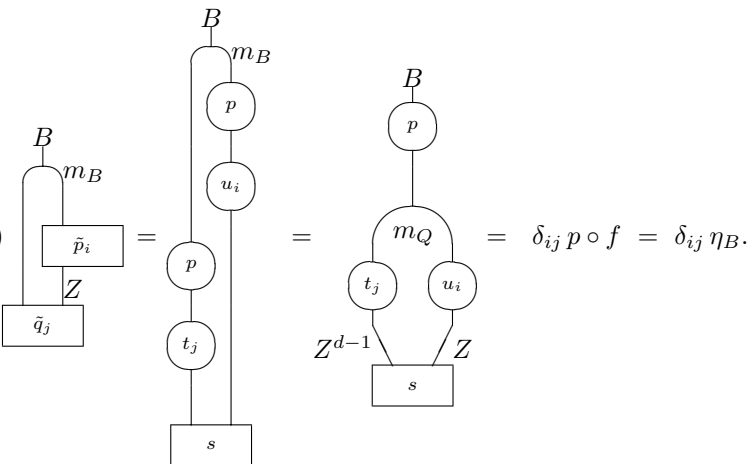
$$(66) \quad p \circ m_Q = m_B \circ p \otimes p,$$

$$(67) \quad p \circ f = p \circ \eta_Q = \eta_B.$$

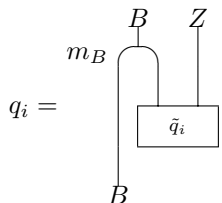
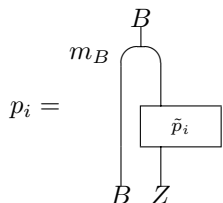
In order prove the claimed isomorphism $B \otimes Z \cong d B$ of B -modules we define morphisms $\tilde{q}_i \in \text{Hom}(\mathbf{1}, B \otimes Z)$, $\tilde{p}_i \in \text{Hom}(Z, B)$, $i = 1, \dots, d$ as the following compositions:

$$\begin{aligned} \tilde{q}_i &: \mathbf{1} \xrightarrow{s} Z^{\otimes d} \cong Z^{\otimes(d-1)} \otimes Z \xrightarrow{t_i \otimes \text{id}_Z} Q \otimes Z \xrightarrow{p \otimes \text{id}_Z} B \otimes Z, \\ \tilde{p}_i &: Z \xrightarrow{u_i} Q \xrightarrow{p} B. \end{aligned}$$

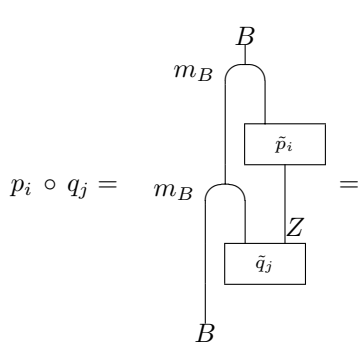
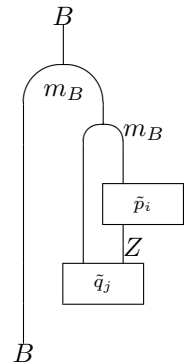
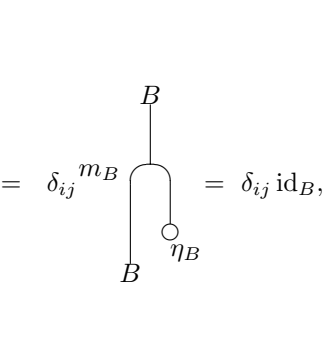
Using, first (66), then (65) and (67) we compute

(68) 

Defining, for $i = 1, \dots, d$,

$q_i =$  $p_i =$ 

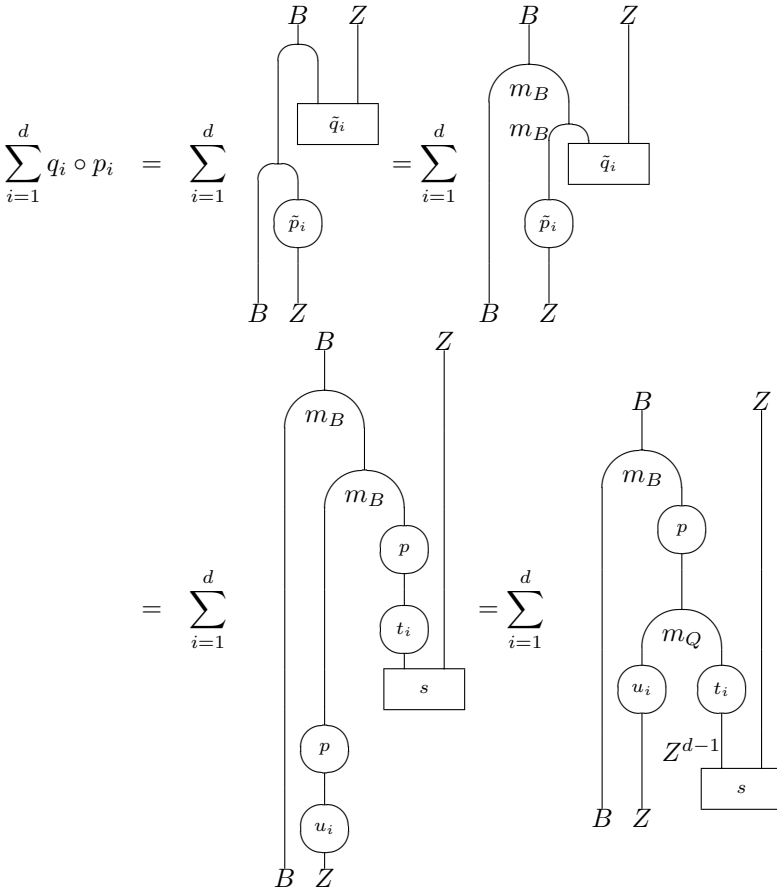
we find

$p_i \circ q_j =$  $=$  $= \delta_{ij} m_B$  $= \delta_{ij} \text{id}_B$,

where in the next to last step we used (68). It is obvious from their definitions that p_i, q_i are morphisms of B -modules. We have thus shown that the B -module $(B \otimes Z, m_B \otimes \text{id}_Z)$ has d direct summands (B, m_B) , and therefore

$$(B \otimes Z, m_B \otimes \text{id}_Z) \cong \underbrace{(B, m_B) \oplus \dots \oplus (B, m_B)}_{d \text{ summands}} \oplus (N, \mu_N).$$

It remains to be shown that $N = 0$ or, equivalently, $\sum_{i=1}^d q_i \circ p_i = \text{id}_{B \otimes Z}$. A short argument to this effect is given in [Deligne, 1990; Bichon, 1998], but since it is somewhat abstract we give a pedestrian computational proof. We calculate



Composition with $\eta_B \otimes \text{id}_Z$ shows that this equals $\text{id}_{B \otimes Z}$ iff

$$(69) \quad \sum_{i=1}^d \text{Diagram} = \text{Diagram}$$

In view of the definition of (Q, m_Q, η_Q) , the left hand side of (69) equals

$$(70) \quad \sum_{i=1}^d (-1)^{d-i} (p \circ c_{S(Z), S(Z)^{\otimes(i-1)}} \otimes \text{id}_{S(Z)^{\otimes(d-i)}} \circ v_1 \otimes \cdots \otimes v_1) \otimes \text{id}_Z \circ \text{id}_Z \otimes s = (p \circ v_1 \otimes \cdots \otimes v_1) \otimes \text{id}_Z \circ \left(\sum_{i=1}^d (-1)^{d-i} c_{Z, Z^{\otimes(i-1)}} \otimes \text{id}_{Z^{\otimes(d-i)}} \otimes \text{id}_Z \circ \text{id}_Z \otimes s \right).$$

Writing $K_i = c_{Z, Z^{\otimes(i-1)}} \otimes \text{id}_{Z^{\otimes(d-i+1)}} \circ \text{id}_Z \otimes s$, where $i \in \{1, \dots, d\}$, one easily verifies

$$\Pi_{d+1}^Z(\sigma_j) \circ K_i = \begin{cases} K_{i-1} & : j = i - 1 \\ K_{i+1} & : j = i \\ -K_i & : \text{otherwise} \end{cases}$$

for all $j \in \{1, \dots, i - 1\}$. This implies that the morphism $Z \rightarrow Z^{\otimes(d+1)}$ in the large brackets of (70) is totally antisymmetric w.r.t. the first d legs, i.e. changes its sign upon multiplication with $\Pi_{d+1}^Z(\sigma_j)$, $j = 1, \dots, d - 1$ from the left. We can thus insert $A_d^Z = s \circ s'$ at the appropriate place and see that (70) equals

$$\begin{aligned} &= (p \circ v_1 \otimes \cdots \otimes v_1) \otimes \text{id}_Z \circ (s \circ s') \otimes \text{id}_Z \\ &\quad \circ \left(\sum_{i=1}^d (-1)^{d-i} c_{Z, Z^{\otimes(i-1)}} \otimes \text{id}_{Z^{\otimes(d-i)}} \otimes \text{id}_Z \circ \text{id}_Z \otimes s \right) \\ &= (p \circ v_1 \otimes \cdots \otimes v_1 \circ s) \otimes \text{id}_Z \\ &\quad \circ \left(\sum_{i=1}^d (-1)^{d-i} s' \otimes \text{id}_Z \circ c_{Z, Z^{\otimes(i-1)}} \otimes \text{id}_{Z^{\otimes(d-i)}} \otimes \text{id}_Z \circ \text{id}_Z \otimes s \right). \end{aligned}$$

Now, $p \circ v_1 \otimes \cdots \otimes v_1 \circ s = p \circ f = \eta_B$. On the other hand, by the total antisymmetry of s we have $s' \circ c_{Z, Z^{\otimes(i-1)}} \otimes \text{id}_{Z^{\otimes(d-i)}} = (-1)^{i-1} s'$ and thus

$$\begin{aligned} \sum_{i=1}^d (-1)^{d-i} s' \otimes \text{id}_Z \circ c_{Z, Z^{\otimes(i-1)}} \otimes \text{id}_{Z^{\otimes(d-i)}} \otimes \text{id}_Z \circ \text{id}_Z \otimes s &= \sum_{i=1}^d (-1)^{d-i} (-1)^{i-1} s' \otimes \text{id}_Z \circ \text{id}_Z \otimes s \\ &= d(-1)^{d-1} s' \otimes \text{id}_Z \circ \text{id}_Z \otimes s = \text{id}_Z, \end{aligned}$$

where the last equality is provided by Lemma 424. Thus (69) is true, implying $\sum_{i=1}^d q_i \circ p_i = \text{id}_{B \otimes Z}$ and therefore the claimed isomorphism $B \otimes Z \cong d(Z)B$ of B -modules. ■

LEMMA 431. *Let \mathcal{C}, Z and the monoid (Q, m, η) be as in Lemma 428. Then the commutative algebra $\Gamma_Q = \text{Hom}(\mathbf{1}, Q)$ is \mathbb{Z}_+ -graded and has at most countable dimension.*

Proof. By construction of Q we have

$$\Gamma_Q = \text{Hom}(\mathbf{1}, Q) = \lim_{\substack{\longrightarrow \\ n}} \bigoplus_{i=0}^n \text{Hom}(\mathbf{1}, S_i(Z)) = \bigoplus_{i \geq 0} \text{Hom}(\mathbf{1}, S_i(Z)).$$

Each of the direct summands on the right hand side lives in \mathcal{C} and thus has finite dimension. It follows that Γ_Q has at most countable dimension. That Γ_Q is a \mathbb{Z}_+ -graded algebra is evident from the definition of m_Q in terms of the morphisms $m_{i,j} : S_i(X) \otimes S_j(X) \rightarrow S_{i+j}(X)$ of Lemma 426. ■

THEOREM 432. *Let $Z \in \mathcal{C}$ be such that $\det Z \cong \mathbf{1}$. Then there exists a commutative monoid (B, m_B, η_B) in $\text{Ind } \mathcal{C}$ such that $\dim \text{Hom}_{\text{Ind } \mathcal{C}}(\mathbf{1}, B) = 1$ and there is an isomorphism $B \otimes Z \cong d(Z)B$ of B -modules.*

Proof. Let (Q, m, η) and the ideal $j = \text{im } m_0 : (J, \mu) \rightarrow (Q, m)$ as before. Assume that j is an isomorphism, thus epi. Then m_0 is epi and thus an isomorphism by Lemma 366. In particular, the map $\Gamma_Q \rightarrow \Gamma_Q$ given by $s \mapsto s \bullet (f - \eta)$ is an isomorphism, thus $f - \eta \in \Gamma_Q$ is invertible. This, however, is impossible since Γ_Q is \mathbb{Z}_+ -graded and $f - \eta \in \Gamma_Q$ is not in the degree-zero part. Thus the ideal j is proper. By Lemma 362 there exists a maximal ideal $\tilde{j} : (\tilde{J}, \tilde{\mu}) \rightarrow (Q, m)$ containing $j : (J, \mu) \rightarrow (Q, m)$. If the monoid (B, m_B, η_B) is the quotient of (Q, m, η_Q) by $\tilde{j} : (\tilde{J}, \tilde{\mu}) \rightarrow (Q, m)$, Proposition 430 implies the isomorphism $B \otimes Z \cong d(Z) \cdot B$ of B -modules. By Lemma 364, the quotient module (B, m_B, η_B) has no proper non-zero ideals, thus by Lemma 365, the commutative \mathbb{C} -algebra $\text{End}_B((B, m_B))$ is a field extending k . By Lemma 358, $\text{End}_B((B, m)) \cong \text{Hom}(\mathbf{1}, B) =: \Gamma_B$ as a \mathbb{C} -algebra. By Lemma 371, the unit $\mathbf{1} \in \text{Ind } \mathcal{C}$ is projective, thus Lemma 363 implies that Γ_B is a quotient of Γ_Q , and by Lemma 431 it has at most countable dimension. Now Lemma 433 below applies and gives $\Gamma_B = \mathbb{C}$ and therefore $\dim \text{Hom}(\mathbf{1}, B) = 1$ as desired. ■

LEMMA 433. *Let $K \supset \mathbb{C}$ a field extension of \mathbb{C} . If $[K : \mathbb{C}] \equiv \dim_{\mathbb{C}} K$ is at most countable then $K = \mathbb{C}$.*

Proof. Assume that $x \in K$ is transcendental over \mathbb{C} . We claim that the set $\{\frac{1}{x+a} \mid a \in \mathbb{C}\} \subset K$ is linearly independent over \mathbb{C} : Assume that $\sum_{i=1}^N \frac{b_i}{x+a_i} = 0$, where the a_i are pairwise different and $b_i \in \mathbb{C}$. Multiplying with $\prod_i (x+a_i)$ (which is non-zero in K) we obtain the polynomial equation $\sum_{i=1}^N b_i \prod_{j \neq i} (x+a_j) = 0 =$

$\sum_{k=0}^{N-1} c_k x^k$ for x . Since x is transcendental, we have $c_k = 0$ for all $k = 0, \dots, N-1$. This gives us N linear equations $\sum_{i=1}^N M_{ki} b_i = 0$, $k = 1, \dots, N$, where $M_{ki} = \sum_{\substack{S \subset \{1, \dots, N\} - \{i\} \\ \#S = k-1}} \prod_{s \in S} a_s$. This matrix can be transformed into the matrix $(V_{ki} = a_i^{k-1})$ by elementary row transformations. By Vandermonde's formula, $\det V = \prod_{i < j} (a_j - a_i) \neq 0$, thus the only solution of $M\mathbf{b} = 0$ is $b_1 = \dots = b_N = 0$, proving linear independence. Since \mathbb{C} is uncountable this contradicts the assumption that K has countable dimension over \mathbb{C} . Thus K/\mathbb{C} is algebraic and therefore $K = \mathbb{C}$ since \mathbb{C} is algebraically closed. ■

Finally we have:

Proof of Theorem 411. If \mathcal{C} is an even STC^* with monoidal generator Z , Lemma 423 allows us to assume $\det Z \cong \mathbf{1}$ (replacing Z by $Z \oplus \bar{Z}$). Now Theorem 432 provides a monoid (B, m, η) in $\text{Ind } \mathcal{C}$ satisfying the assumptions of Corollary 415, which gives rise to a symmetric fiber functor $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$. ■

REMARK 434. It seems instructive to point out the main difference of our proof of Theorem 411 w.r.t. the approaches of [Deligne, 1990; Bichon, 1998]. In [Deligne, 1990], a commutative monoid (Q, m, η) for which there is an isomorphism $Q \otimes Z \cong d(Z)Q$ of Q -modules is constructed by a somewhat complicated inductive procedure. The explicit construction of the monoid that we gave is due to [Bichon, 1998]. Deligne proceeds by observing that, for every $X \in \mathcal{C}$, the k -vector space $\text{Hom}(\mathbf{1}, Q \otimes X)$ is a module over the commutative ring $\Gamma_Q := \text{End}_Q((Q, m)) \cong \text{Hom}(\mathbf{1}, Q)$, and the functor $\tilde{E} : X \mapsto \text{Hom}(\mathbf{1}, Q \otimes X)$ is monoidal w.r.t. the tensor product of $\Gamma_Q - \text{Mod}$ (rather than that of $\text{Vect}_{\mathbb{C}}$). Now, a quotienting procedure w.r.t. a maximal ideal J in Γ_Q is used to obtain a tensor functor $E : \mathcal{C} \rightarrow K - \text{Vect}$, where $K = \Gamma_Q/J$ is a field extension of the ground field k . If $\text{Hom}(\mathbf{1}, Q)$ is of at most countable dimension then $[K : k] \leq \aleph_0$, and if k is uncountable and algebraically closed it follows that $K = k$.

Our approach differs in two respects. Less importantly, our insistence on $\det Z \cong \mathbf{1}$ makes the construction of the monoid (Q, m, η) slightly more transparent than in [Bichon, 1998]. More importantly, we perform the quotienting by a maximal ideal inside the category of Q -modules in $\text{Ind } \mathcal{C}$ rather than in the category of Γ_Q -modules, yielding a monoid (Q', m', η') in $\text{Ind } \mathcal{C}$ with $\Gamma_{Q'} = \mathbb{C}$. Besides giving rise to a symmetric fiber functor $E : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{C}}$ in a more direct fashion, this has the added benefit, as we will show in the final subsection, of allowing to recover the group $\text{Nat}_{\otimes} E$ without any reference to the fiber functor and its natural transformations! The ultimate reason for this is that, due to uniqueness of the embedding functor, the monoid (Q', m', η') in $\text{Ind } \mathcal{C}$ is nothing but the monoid $(\pi_l, \tilde{m}, \tilde{\eta})$ in $\text{Rep } G$ that arises from the left regular representation of G , cf. [Mürger and Tuset, 2006].

B.12 Addendum

In the previous subsection we have concluded the proof of the existence of a fiber functor and, by the concrete Tannaka theorem, of the equivalence $\mathcal{C} \simeq \text{Rep}_f(G, k)$, where (G, k) is a compact supergroup. However, we would like to show how the group $\text{Nat}_{\otimes} E$, and in some cases also G , can be read off directly from the monoid (Q, m, η) , bypassing fiber functors, natural transformations etc.

DEFINITION 435. The automorphism group of a monoid (Q, m, η) in a strict tensor category \mathcal{C} is

$$\text{Aut}(Q, m, \eta) = \{g \in \text{Aut } Q \mid g \circ m = m \circ g \otimes g, g \circ \eta = \eta\}.$$

PROPOSITION 436. Let \mathcal{C} be an STC* and (Q, m, η) a monoid in $\text{Ind } \mathcal{C}$ satisfying

(i) $\dim \text{Hom}_{\text{Ind } \mathcal{C}}(\mathbf{1}, Q) = 1.$

(ii) For every $X \in \mathcal{C}$, there is $n(X) \in \mathbb{Z}_+$ such that $n(X) \neq 0$ whenever $X \not\cong 0$ and an isomorphism $\alpha_X : (Q \otimes X, m \otimes \text{id}_X) \rightarrow n(X) \cdot (Q, m)$ of Q -modules.

Then the group $\text{Nat}_{\otimes} E$ of monoidal natural automorphisms of the functor constructed in Proposition 413 is canonically isomorphic to the group $\text{Aut}(Q, m, \eta)$.

Proof. Let $g \in \text{Aut}(Q, m, \eta)$. For every $X \in \mathcal{C}$ define $g_X \in \text{End } E(X)$ by

$$g_X \psi = g \otimes \text{id}_X \circ \psi \quad \forall \psi \in E(X) = \text{Hom}(\mathbf{1}, Q \otimes X).$$

From the definition of $(g_X)_{X \in \mathcal{C}}$ and of the functor E it is immediate that $(g_X)_{X \in \mathcal{C}}$ is a natural transformation from E to itself. We must show this natural transformation is monoidal, i.e.

$$\begin{array}{ccc} E(X) \otimes E(Y) & \xrightarrow{d_{X,Y}} & E(X \otimes Y) \\ g_X \otimes g_Y \downarrow & & \downarrow g_{X \otimes Y} \\ E(X) \otimes E(Y) & \xrightarrow{d_{X,Y}} & E(X \otimes Y) \end{array}$$

commutes. To this end consider $\phi \in E(X) = \text{Hom}(\mathbf{1}, Q \otimes X), \psi \in E(Y) = \text{Hom}(\mathbf{1}, Q \otimes Y)$ and $g \in \text{Aut}(Q, m, \eta)$ with $(g_X)_{X \in \mathcal{C}}$ as just defined. Then the image of $\phi \boxtimes \psi \in E(X) \otimes E(Y)$ under $g_{X \otimes Y} \circ d_{X,Y}$ is

$$g \otimes \text{id}_{X \otimes Y} \circ m \otimes \text{id}_{X \otimes Y} \circ \text{id}_Q \otimes \phi \otimes \text{id}_Y \circ \psi,$$

whereas its image under $d_{X,Y} \circ g_X \otimes g_Y$ is

$$m \otimes \text{id}_{X \otimes Y} \circ g \otimes g \otimes \text{id}_{X \otimes Y} \circ \text{id}_Q \otimes \phi \otimes \text{id}_Y \circ \psi.$$

In view of $g \circ m = m \circ g \otimes g$, these two expressions coincide, thus $(g_X) \in \text{Nat}_{\otimes} E$. It is very easy to see that the map $\sigma : \text{Aut}(Q, m, \eta) \rightarrow \text{Nat}_{\otimes} E$ thus obtained is a group homomorphism.

We claim that σ is an isomorphism. Here it is important that we work in $\text{Ind } \mathcal{C}$ rather than any category $\widehat{\mathcal{C}}$, since this implies that Q is an inductive limit of objects in \mathcal{C} . The assumptions (i),(ii) then give $\text{Hom}(X, Q) \cong \text{Hom}(\mathbf{1}, Q \otimes \overline{X}) \cong \mathbb{C}^{n(\overline{X})}$ for all $X \in \mathcal{C}$ and thus (using $n(X) = n(\overline{X}) = \dim E(X)$)

$$(71) \quad Q \cong \varinjlim_{s \subset I} \bigoplus_{i \in S} n(X_i)X_i \quad \text{and} \quad \text{End } Q \cong \prod_{i \in I} \text{End } E(X_i),$$

where S runs through the finite subsets of I . Assume now that $\sigma(g)$ is the identity natural transformation, i.e. $g \otimes \text{id}_X \circ \phi = \phi$ for all $X \in \mathcal{C}$ and $\phi \in \text{Hom}(\mathbf{1}, Q \otimes X)$. Be the existence of conjugates in \mathcal{C} , this is equivalent to $g \circ s = s$ for all $Y \in \mathcal{C}$ and $s \in \text{Hom}(Y, Q)$. Since Q is an inductive limit of objects in \mathcal{C} , this implies $g = \text{id}_Q$.

If now $\alpha \in \text{Nat}_{\otimes} E$, we first observe that α is a natural isomorphism by 396. By the isomorphisms $\text{Nat } E \cong \prod_{i \in I} \text{End } E(X_i)$ (cf. the proof of Proposition 375) and (71), we have a map $\text{Nat}_{\otimes} E \rightarrow \text{Aut } Q$. Reversing the preceding computations shows that every $\alpha \in \text{Nat}_{\otimes} E$ gives rise to an element of $\text{Aut}(Q, m, \eta)$. ■

REMARK 437. This result shows that the group $\text{Nat}_{\otimes} E$ can be recovered directly from the absorbing monoid (Q, m, η) in $\text{Ind } \mathcal{C}$. In general the compact group G as defined in Subsection B.1 is a true subgroup of $\text{Nat}_{\otimes} E$, the latter being the pro-algebraic envelope of G . (In the cases of $G = U(1), SU(2), U(2)$, e.g., that would be $\mathbb{C}^\times, SL(2, \mathbb{C}), GL(2, \mathbb{C})$, respectively.) But if \mathcal{C} is finite (i.e. has finitely many isomorphism classes of simple objects) then $\text{Nat}_{\otimes} E$ is finite and $G = \text{Nat}_{\otimes} E$. Interestingly, even in the case of finite \mathcal{C} , where the monoid (Q, m, η) actually lives in \mathcal{C} , there seems to be no way to recover G without using $\text{Ind } \mathcal{C}$ at an intermediate stage.

INDEX

- ABL-rule, 597
absolute objects, 1343
absolutism, 18
abundance of elements, 1221
accelerating expansion, 1209
adiabatic measurement, *see* measurement, weak
adjoint representation, 68, 69
age of the universe, 1196, 1229
Aharonov, Y., 596, 597
Albert, D., 597
Anderson, J., 1343, 1348, 1350, 1353
angular power spectrum, 1202
anomalies, 711
apparent size, 1197, 1199
area distance, 1197
Aristotle, 1313
astronomical observation, 1220, 1221
asymptotic freedom, 714
averaging scale, 1223
- background independence, 1307
baryosynthesis, 1193, 1232
BB84 (Bennett and Brassard protocol), 595, 596
BBGKY approach, 1034–1038
BCS theory of superconductivity, 493, 1122–1125
Beethoven, L., 529
beginning of the universe, 1212
Bekenstein entropy, 1294, *see* Hawking effect
Bell locality, 386, 387, 391, 395
Bell states, 567, 591
Bell's inequality, 386, 566, 596, 617, 638, 754
Bell's theorem, 308, 384, 386, 396, 755
- Bell-type experiment, 396
Belot, G., 209
Benioff, P., 556
Bennett, C., 556, 590, 593, 594, 601
Berezin function, 480
 classical limit, 480
 quantization, 459
Bergmann, P., 596, 1291
Bernoulli, D., 943
best-fit parameters, 1227
big bang, 209, 491, 1169, 1185, 1191–1194, 1209, 1210, 1212, 1213, 1214, 1236, 1287, 1318, 1322, 1424, 1425, *see* initial singularity
biorthogonal decomposition, 288, 358, 566, 578, 607, 611
bit, 558
bit commitment, 594, 600
 binding protocol, 606
 concealing protocol, 606
 'no go' theorem, 594, 606
 unconditionally secure, 601, 637, 639
black hole, 1086–1088, 1287, 1294
 thermodynamics, 1295, *see* Hawking effect
blackbody spectrum, 1077, 1197
Bloch sphere, 295, 296, 309
Bohm, D., 323, 385, 418, 437, 566, 653
Bohm interpretation of quantum mechanics, 373–377, 392, 646, 651–653, 762, 1404, 1406, *see* hidden variables
Bohr, N., 424, 433–444, 555, 647
Bohr's atomic model, 424, 483
Bohr's correspondence principle, 426,

- 429
- Boltzmann, L., 952–992
- Boltzmann equation, 962–969, *see*
 BBGKY approach
 Lanford’s approach, 1028–1033
- Boltzmann’s combinatorial argument
 and responses, 975–983, 983–
 992
- Bondi, H., 650
- Boolean algebra, 642
- Born rule, 524
- Bose–Einstein
 condensation, 1082–1084
 statistics, 848–855
- Bose gas, 1125
- bra-ket notation, *see* Dirac bra-ket
 notation
- braided tensor category (BTC*), 870
- braiding, 867
- Brassard, G., 556, 590, 594, 601, 650
- Briegel, H., 616
- broadcasting, 580
 ‘no broadcasting’ theorem, 581
- Broglie-Bohm theory, *see* Bohm in-
 terpretation of quantum me-
 chanics
- Bronstein, M., 1291
- Brout–Englert–Higgs mechanism, 693–
 698
- Brukner, Č., 632
- BTC* (braided tensor category), 870
- Bub, J., 632, 646
- Bub–Clifton theorem, 646
- Butterfield, J., 138, 155, 209, 397
- C^* -algebra, 419, 445, 451–463, 468–
 474, 492–514, 520–522, 633–
 637, 651, 653, 732–740, 751,
 754, 755, 841, 843–845, 894,
 1110–1113, 1120, 1144, 1151,
 1163, *see* GNS construction,
 von Neumann algebra
- canonical action, 103
- canonical quantization, 447
- Carnot’s theorem, 934
- category, 863
 *-category, 869
 C^* -, 874
 \mathbb{C} -linear, 869
 Ab-, 868
 Abelian, 875
 additive, 869
 braided monoidal, 867
 equivalence, 864
 complex-linear, 869
 monoidal, 865
 of localized transportable endo-
 morphisms, 783f
 product, 865
 symmetric monoidal, 867
- Cauchy
 development, 1415–1418
 horizon, 1416–1417
 surface, 1394, 1409–1411, *see* global
 hyperbolicity
- causal properties of relativistic space-
 times, 263
 past distinguishing, 271
 chronology, 271
 closed timelike curves, 271, 1421
 future distinguishing, 271
 global hyperbolicity, 1410, 1419,
 1423–1424, *see* Cauchy sur-
 face
 past distinguishing, 271
 stable causality, 1368, *see* global
 time function
 strong causality, 271, 1410
- CBH (Clifton–Bub–Halvorson) theo-
 rem, 632, 650
- CBR (cosmic background radiation),
 1192, 1195, 1197, 1200, 1202–
 1206, 1215, 1223, 1226, 1231
 anisotropies, 1209, 1214, 1223, 1224,
 1227, 1231
 isotropy, 1195, 1225, 1226
 polarization, 1210
 power spectrum, 1202, 1203, 1217

- temperature, 1229
- CCR (canonical commutation relations), 287, 334, 340–341, 447, 1125, 1423
- Chalmers, A.,
 - on Curie’s principle, 1334
- channel capacity, 562, 564
- chaos
 - classical, 490, 491, 520, 1390, 1403
 - deterministic, 1390, 1391
 - quantum, 478, 486–491, 519, 1106, 1402
 - intrinsic, 490
- Chau, H., 594, 601
- chronology, *see* causal properties of relativistic spacetimes
- chronology protection conjecture, 1419
- Church, A., 614
- Church–Turing thesis, 614, 1420, *see* Turing machine
- classical limit (of quantum mechanics), 360, 471–530, *see* high temperature limit, thermodynamic limit, van Hove limit
 - the limit $\hbar \rightarrow 0$, 471–491, 1156–1157
 - the limit $N \rightarrow 0$, 492–514
- classical observables, 493, 512, 514ff, 1162–1163
- classical probabilities, 283, 285, 309
- classical states, 316, 422, 472, 500, 514
- classical world, 420
 - absolute existence of, 420
- Clausius, R., 932–938, 944
- Clifton, R., 632, 646
- clock (or twin) paradox, 236
- cloning, 582
 - ‘no cloning’ theorem, 580, 581
 - and measurement, 652
- closed timelike curves, 271, 1419 *see* causal properties of relativistic spacetimes 1419
- clustering, 740, 1108, 1109
- CNOT gate, 616, 632
- co-adjoint representation, 73
 - symplectic structure of, 99
- coarse-graining (for consistent histories), 526
- coherent state, 429, 455, 472, 519
 - Schrödinger’s, 455
- cokernel, 875
- collapse interpretations (of quantum measurement), 377–382, 645–647
- collapse postulate, 363, 377, 378, 397, 1165
- common cause, 382–386
- complementarity, 441, 555, 647
 - early Bohr, 441
 - Heisenberg, 442
 - Landsman, 443
 - later Bohr, 442
 - Pauli, 443
- completely positive map, 456, 574, 755
- completeness condition, 570
- complexity class
 - EXP**, 614
 - NP**, 615
 - P**, 614
- compression/decompression scheme, 561, 589
- computational complexity, 614
 - easy (tractable), 614
 - hard (intractable), 614, *see* complexity class
- concordance model, 1214–1216
- Connes, A., 1295
- confinement, 721
- conservation of energy-momentum (in general relativity), 240–243
- conservation principles (in general relativity), 252–256, *see* Noether’s theorems
- consistent histories, 419, 422, 517, 518, 522–529
- constant mean curvature time, 214–

- 221
- constraint equation, 1230, 1352, 1372, 1412, 1424
- constraints (in Hamiltonian mechanics), 151–153, 180
- first class, 180, 1379
- primary, 1379
- secondary, 1379
- contact transformations, 1341
- container view of spacetime, 18, 1312, 1377, 1411
- continuous field of C^* -algebras, 457
- continuous spontaneous localisation, 380
- Copenhagen interpretation, 349, 422, 433–446, 555, 645, 655, 1257, 1232
- Copernican Principle, 1225, 1226, 1273
- correspondence principle, 418, 424–429, 483
- cosmic censorship conjecture, 1419, 1425
- cosmic variance, 1203, 1217
- cosmological constant, 243–249, 1187, 1188, 1190, 1197, 1200, 1202, 1227, 1408, *see* Einstein's field equations
- cosmological principle, 1223, 1231
- cosmological time function, 216
- cotangent bundle
- symplectic structure of, 9, 122
- cotangent lift, 18, 20, 58, 116
- covariance, 330, 335, 1343
- general, *see* general covariance
- Crépeau, C., 590
- creation of the universe, 1212, 1218, 1230, 1236, 1271
- critical density, 1188, 1199, 1201
- Curie's principle, 1333–1336, 1346, 1374–1375
- Curie, P., 1374–1375, 1415
- cyclic (ignorable) coordinates, 5, 15
- Darboux's theorem, 10, 85, 91, 97, 98, 142
- dark energy, 1185, 1188, 1200, 1202, 1216, 1224, 1227, 1232, 1267, 1268, 1408
- dark matter, 1201
- cold dark matter (CDM), 1187, 1189, 1201, 1202, 1244
- hot dark matter, 1201
- de Broglie-Bohm theory, *see* Bohm interpretation of quantum mechanics
- De Finetti's representation theorem, 500
- de Sitter spacetime, 213, 215, 1213,
- de Sitter universe, 1189, 1190, 1198, 1204, 1212
- Debye, P., 1080
- deceleration parameter, 1187, 1197, 1198, 1227
- decoherence, 211, 365–368, 372, 373, 515–529, 646, 654, 1401, 1405
- deformation quantization, 456
- functorial reinterpretation, 470
- delay-differential equations, 1371
- Deligne, P., 901, 919
- density operator, 281–285, 288, 289, 295, 296, 302, 307, 313, 317, 376
- density parameter, 1187, 1188, 1227
- Descartes, R., 1312–1313
- Deser, S., 1292
- determinism
- as a methodological imperative, 1372
- varieties of, 1373–1374
- Laplacian, 1370
- non-Laplacian, 1373–1374
- Deutsch, D., 556, 617, 631
- Deutsch's XOR algorithm, 617, 618
- Deutsch–Jozsa algorithm, 618, 621
- Dieks, D., 580
- diffeomorphism, 31
- diffeomorphism invariance, 1308, *see* general covariance

- differential equation
 - ordinary, 32, 1373, 1385
 - partial
 - elliptic, 1382
 - hyperbolic, 1395, 1396
 - parabolic, 1382
 - weak solutions, 1384
- dimension in a category, 871
 - integrality of, 907
- Dirac bra-ket notation, 296–198, 305
- Dirac equation, 677–679, 720
- Dirac matrices, 469, 678, 687
- Dirac operator, 471, 1295, 1296
- Dirac picture, 306, 1161
- Dirac, P. A. M., 153, 377, 378, 427, 429, 432, 447, 459, 462, 467, 469, 470, 473, 555, 642, 718, 1239, 1291, 1379
- disjoint representations, 498, 732, 843, 844
- disjoint states, 498, 843
- distances of observed objects, 1195
- domain of causal influence, 1205
- domain of dependence, 241, 1387m
 - see* Cauchy development
- domain of prediction, 1397–1398, 1419
- Doplicher, S., 744, 766, 800, 801, 815, 837, 840, 845, 855
- Doplicher–Haag–Roberts (DHR)
 - theory of superselection, 502, 781ff
- double slit experiment, 345, 377, 441, 443, 524, 536
- dual of a Lie algebra, 88, 90
- dynamical objects, 1343
- dynamical reduction theories, 379
- dynamical triangulation, 1288, 1296
- Earman, J., 177, 197, 206, 209
 - on Curie’s principle, 1335
- Eddington, A. S., 514, 1088
- EGS theorem, 1226
- Ehrenfest time, 478
- Ehrenfest’s theorem, 475
- Ehrenfests, on Boltzmann, 930, 980, 1005
- eigenstate–eigenvalue link, 285, 286, 357–359, 367, 378, 380, 381, 1405–1406
- einselection, 516–522, *see* decoherence
- Einstein, A., 206, 252, 258, 259, 348, 350, 418–424, 427, 433, 435, 442, 471, 478, 486, 488, 555–557, 639, 640, 645, 647–652, 940, 941, 977, 989, 991, 993, 1003, 1021, 1075–1082, 1290, 1309–1314, 1346–1353, 1356, 1357, 1363, 1410–1414
- Einstein Field Equations (EFE), 200, 243–245, 249, 263, 1185, 1187, 1190, 1224, 1230, 1290, 1352, 1409–1416
- Einstein’s ‘hole argument’, 1308, 1349, 1351, 1411
- Einstein’s principle of relativity, 1346
- Einstein static universe, 1184, 1189, 1212
- Einstein–de Sitter model (or universe), 1189, 1198, 1204
- Ekert, A., 596, 599
- embedding, 41
- empirical content of quantum theory, 307–308, 321–339
- energy-density conservation equation, 1186
- energy conditions (in general relativity), 240, 1207
 - dominant, 240, 1394
 - strong, 1187
- ensemble interpretations,
 - of quantum mechanics, 363, 432, 478
 - of statistical mechanics, 943
- entangled state (in quantum mechanics), 288, 367, 368, 396, 445, 565, 566, 754–756
- entanglement (quantum), 397, 439, 444–446, 518, 555, 556, 590–

- 594, 602, 603, 605, 611, 612, 616, 630, 635–640, 647, 649–651, 754–756
- entropy
 conditional, 562, 563, 575, 577
 gravitational, 1208
 joint, 563, 577, 578
 relative, 564
 Shannon, 557, 558, 562, 564, 576, 577
 von Neumann, 575, 577, 579, 1091
- environment-induced superselection, 516–522, *see* decoherence
- EPR (Einstein–Podolsky–Rosen) argument (or paradox), 308, 349–355, 385, 555, 556, 590, 639
 EPR criterion for physical reality, 351–353, 384
 EPR attack (or cheating strategy), 601, 603, 605, 606, 609, 610, 637
- equivariance, 63, 68, 75
 of momentum maps, 103, 112–114, 117
- ergodic theory (classical) 1005–1020, 1025, *see* quantum ergodic theory
 ergodic hierarchy, 1011–1020, 1104–1105, 1390
 ergodic hypothesis in early Boltzmann, 955–962, *see* Boltzmann’s combinatorial argument *and* H-theorem
 ergodic theorem, 1006–1009, 1093
 Khinchin’s approach, 1017–1028
- Euclidean group, 17–20, 62, 66, 103, 1100, 1113, 1151, 1154
- Euler
 equations for rigid body, 81, 91
 equations for fluid, 16, 1383–1384
- Euler–Lagrange equations, 145, 146, 158, 159, 184, 190, 208, 688, 1356–1358, 1378, 1379
- Everett interpretation, 370, 646, 1407–1408
 Saunders-Wallace version, 646
- evolution of the universe, 1185, 1188, 1233, 1267, 1273, 1319
- existential interpretation (of quantum mechanics), 517
- expansion of the universe, 1183, 1184, 1187, 1192, 1196, 1200, 1214, 1221, 1239, 1245, 1252,
- Faddeev, L., 689–691, 695, 1292
- Faraday lines, 1300
- Faraday, M., 1300, 1314, 1324
- Fermi–Dirac statistics, 679, 704, 719, 845–855, 1085
- Feynman, R., 556, 1170, 1291, 1292, 1391
- Feynman diagram, 664–667, 676, 682, 699, 704, 709
- Feynman path integral, 671–674
- Feynman propagator, 666, 698
- Feynman rules
 classical theory, 663–669
 quantized theory, 674–677
 Dirac fields, 681–683
 calculation of scattering matrix, 690–691
 perturbative quantum general relativity 1292
- fiber functor, 881
 existence, 884, 900, 902
 uniqueness, 882, 893
- fidelity, 561, 588, 589
- Fierz, M., 1290
- fine-graining (for consistent histories), 526
- fine tuning, 723, 1205, 1209, 1212, 1215, 1218, 1238, 1248, 1249, 1264, 1265, 1285
- finitely generated
 algebra, 918
 tensor category, 900
- folium, 498

- Friedman, M., 197, 206
- Friedmann equation, 1187, 1188, 1196, 1198, 1211, 1217
- Friedmann–Lemaître (FL) model (or universe), 1184, 1187, 1188, 1190, 1195, 1204, 1207–1209, 1212, 1215, 1222–1229, 1234, 1238, 1266, 1269, 1270, 1275
- Friedmann–Robertson–Walker model, 1187, 1322, 1426, 1427, *see* Friedmann–Lemaître model
- Frobenius’ theorem, 7, 29, 38–43, 51, 94, 97, 98, 151, 250, *see* Perron–Frobenius’ theorem
- Fuchs, C., 580, 601, 632, 650
- functor, 864
 - *-preserving, 869
 - braided monoidal, 867
 - essentially surjective, 864
 - faithful, 864
 - fiber, 881
 - full, 864
 - monoidal, 866
 - symmetric monoidal, 867
- funnel property, 743
- future distinguishing, *see* causal properties of relativistic spacetimes
- Gödel, K., 614
- Galilean group, 25, 299, 329–334
- Galilean transformation, 299, 337, 649, 1375, 1397, 1398
- Galois, É., 1337–1338
- gauge conditions
 - Lorentz gauge, 1378, 1379
- gauge fixing, 153, 688, 692, 704, 716, 1318, 1378
- gauge freedom, 151–153, 177–191, 1352, 1376, 1378–1380, 1410, 1413
- gauge symmetry, 197, 203, 467, 685, 687, 719, 806, 809, 811, 815, 830, 846–848, 853, 854, 1124, 1355, 1359, 1378–1380, 1410, 1413
- general covariance, 133, 136, 196–219, 1302, 1306, 1310, 1311, 1343, 1348–54, 1357, 1358, 1363, 1409, 1411, 1414
- general linear group, 44, 49, 409
- geometric quantization, 462–467, 471, *see* quantization
 - functorial reinterpretation, 470
- geometry of the universe, 1222
- Ghirardi, G.-C., 379–381, 420, 522, 645, 1403
- GHZ (Greenberger–Horne–Zeilinger) theorem, 361, 387
- Gibbs’ statistical mechanics, 992–1005, *see* BBGKY approach
- Gisin, N., 379, 569
- Gleason’s theorem, 281, 282, 286, 295, 308, 634, 644, 654, 1089, 1405
- global hyperbolicity, 1410, 1421–1424, *see* Cauchy surface
- global observable, 1101, 1109, *see* macroscopic observables
- global time function, 1370, 1410, 1411, 1419
- GNS (Gelfand–Naimark–Segal) construction (or representation), 318, 442, 445, 498, 499, 502, 503, 505, 508, 510, 732, 740, 758, 759, 782, 800, 806, 827, 829, 830–833, 836, 843, 845, 851, 1089, 1101, 1109, 1110–1113, 1117, 1119, 1120, 1121, 1124–1127, 1133, 1144, 1148, 1152, 1154, 1163
- Goroff, M., 1292
- gravitational lensing, 1197–1199, 1202, 1203, 1214, 1227
- gravitational instability, 1193, 1210, 1214
- gravitational waves, 1186, 1203, 1207, 1210, 1232, 1252, 1290, 1291
- GR, *see* GTR (general theory of relativity)

- Green, M., 1288, 1298
 groups and group theory, 299, 409, 685, 1337–1342
 amenable, 740, 1113
 continuous, 44, 1340–1342
 discrete, 1340
 Lie, 7, 8, 15, 19, 20, 29, 37, 40, 44–79, 147, 1341, 1342, 1376
 matrix, 44, 50, 51, 71, 72
 semi-simple, 1113
 compact, 900, 901
- GRW (Ghirardi–Rimini–Webber) ‘collapse theory’, 645, 651
- GTR (general theory of relativity), 133–138, 147, 156, 172, 178, 184, 196–221, 229, 246, 263, 384, 467, 724, 1185, 1230, 1287, 1348–1354, 1409–1423, *see* Einstein field equations
- Gupta, S., 1290
- Gutzwiller trace formula, 487, 488
- H*-theorem, 962–969
 objections and responses, 970–974
- Haag, R., 738ff, 777, 781ff, 1099, 1124, 1125
- Hadamard gate, 616, 632
- Hadamard transformation, 618, 620–622, 624–626
- Hadamard, J., 1105, 1390
- Halvorson, H., 569, 632, 633
- Hamilton’s equations, 11, 85, 87, 90, 141
- Hamilton–Jacobi equations of GR, 1293
- Hamiltonian approach to mechanics and field theory, 8–13, 79–119, 140–143, 147–154, 165–168, 180, 193
- Hamiltonian, 303, 304–307, 323, 325
- Hamiltonian symmetries, 12, 58
- Hamiltonian system, 12, 150
- Hamiltonian vector field, 12, 85, 458
- Hardy’s axioms (for quantum mechanics), 320–322
- Hardy’s no-go theorem, 387–389
- Hartle, J., 518, 522, 528, 530, 546, 1262, 1264, 1294, 1315, 1323
- Hartle–Hawking ‘no-boundary’ proposal, 1213, 1236
- Hausdorff property (of topological spaces), 63, 147, 155, 175, 445, 457, 507, 635, 885, 894, 1295, 1409, 1418
- Hawking effect, 1143, 1144, 1294, 1424, *see* Unruh effect
- Hawking, S., 244, 271, 1143, 1207, 1213, 1236, 1294, 1424
- Healey, R., 191, 209, 375, 421
- heat equation, 927, 1096, 1380–1381
- Heisenberg, W., 340, 342, 417, 418, 420, 421, 426–448, 459, 471, 484, 486, 515, 523, 529, 555, 641, 1126, 1291, 1292
- Heisenberg cut, 418, 437–441, 446, 518
- Heisenberg picture, 306, 307, 448, 475, 520, 521, 523, 635, 1115, *see* Schrödinger picture
- Heisenberg equation, 307, 509
- Heisenberg’s matrix mechanics, *see* matrix mechanics
- Helmholtz, H. von, 159, 339, 1381
- Helmholtz free energy, 939, 1091
- Henneaux, M., 153, 467, 1379–1380
- Hepp, K., 477, 492, 493, 499, 502, 512, 513, 1163
- Hertz, H., 28
- Hertz, P., 995
- hidden variables (in quantum mechanics), 372–376, 384–395, 431, 446, 646, 648, 653, 754, 1401–1406, *see* Bohm interpretation
- Higgs mechanism, *see* Brout–Englert–Higgs mechanism
- high temperature limit, 1078, 1081,

- 1157–1159
- Hilbert, D., 1076, 1268, 1270, 1348, 1351, 1357, 1361, 1363
- Hilbert space, 397–400
 formalism for quantum statistical physics, 1089–1094
 limitations of, 1094–1099
- hole argument, *see* Einstein's hole argument
- Holevo bound, 582, 585, 593
 derivation, 583
- holism, 395
- horizons, *see* physics horizon
 event, 1203, 1417
 Cauchy, 216, 1414, 1415
 causal, 1203, 1204, 1209
 particle, 1203, 1204, 1206, 1209, 1215, 1230
 visual, 1203, 1204, 1206, 1209, 1226, 1227, 1230, 1231, 1247, 1272
- horizon problem, 1204, 1209, 1238
- hot big bang, *see* big bang
- Hubble constant (or parameter), 1186, 1187, 1190, 1196, 1197, 1227
- Hughston–Jozsa–Wootters theorem, 569, 590
- Husimi function, 480
- hydrodynamics as part of statistical physics, 926
- Hyperion, 418, 478, 519, 1403
- ideal fluid, 8, 15, 16, 82
- ignorable coordinates, *see* cyclic coordinates
- ignorance interpretation (of mixed states), 357, 432, 433, 439, 440, 512, 516, 544, 749
- Immirzi parameter, 1295, 1301
- improper mixture, *see* mixed state, improper
- impulsive measurement, 321–324, 327, 337, 377
- incompatible observables, 285, 308, 346–349
- incomplete motion, 491, 492, 1382, 1385, 1400
- inductive limit C^* -algebra, 496, 504, 739, 765, 829, 879, 900, 901, 911, 920, 921, 1099, 1100, 1107, 1129, 1141, 1151, 1152
- inertial frame, 140, 164, 169, 170, 188, 206, 334–339, 350, 451, 648, 649, 1297, 1375, 1376
- inflationary model (or universe), 1209–1212, 1215, 1224, 1233, 1234, 1284
 chaotic, 1206, 1210, 1213, 1219, 1227, 1231, 1234–1236, 1244, 1245, 1256, 1262, 1266, 1269
- information
 accessible, 582
 channel, 562
 classical, 557, 576, 577
 compressibility, 576, 577
 compression, 562, 578
 'block coding', 561
 blind, 586
 classical, 557
 quantum, 585
 visible, 586
 compression rate, 560, 586
 inaccessible, 577
 mutual, 562, 564, 575, 577
 quantum, 565, 576, 577
 represented physically, 576
 Shannon's sense, 575
 source, 557, 561
 ergodic, 558
 stationary, 557
- information loss paradox, 1424, 1426
- initial singularity (of the universe), 1188, 1190, 1204, 1211, 1232, 1235, 1236, 1284, 1426, *see* big bang
- initial conditions for the universe, 1223, 1233, 1237, 1238, 1273, 1426
- initial value formulation/problem, 1370,

- 1381, 1385
- properly posed, 1390
- well posed, 1388–1390, 1395, 1414
- instantons, 719
- integrable system, 482, 483, 485
 - classical limit, 482
- integral manifold, 42, 43, 98, 151
- interaction Hamiltonian, 322, 324, 336, 364
- interaction picture, 306, 1161
- interference pattern, 345–347, 378, 441,
 - see* wave-particle duality
- intermediate isotropisation, 1226
- invariance, 328, 333, 337, 1331, 1343
- invariant means and states, 740, 1108
- involution, 42
- Isham, C., 209, 452, 462, 522–525, 1212, 1254, 1257, 1294, 1313, 1318, 1324
- isotropic universe, 1225, 1226, *see* Copernican Principle
- isotropy group, 65, 66, 68, 120, 128

- Jacobson, T., 1293, 1303
- Jaynes, E., 569, 932, 943, 977, 993, 1003, 1010, 1090
- Jozsa, R., 557, 569, 580, 585, 588, 590, 616, 618, 621, 627, 630–632

- Kaluza–Klein theory, 28, 29
- Kent, A., 517, 522, 527, 602, 608
- Kepler, J., 1305
- Kepler problem, 5, 15, 174–175, 483, 484
- key distribution and storage, 556, 594, 595–596, 599
 - public, 617
 - secure, 602, 650
- Killing fields, 165, 199, 215, 252–256, 1421
- kinematic independence, 633, 635, 636, 639, *see* microcausality
- kinematics on Lie groups, 75
- kinetic theory of gases, 927, 941–952

- Klein, F.
 - and group theory, 1338–1342
 - and his Erlangen Program, 1339–1340
- Klein, M., 649, 930, 931, 957, 962, 965, 974, 990
- Klein, O., 1290
- Klein-Gordon equation, 154, 192, 207, 208, 1394, 1396, 1416, 1422
- KMS (Kubo–Martin–Schwinger) states, 1114–1154
 - extremal KMS states and pure phases, 499, 1144–1154
 - in BCS theory, Bose gas and Heisenberg model, 1122–1128
 - in QFT, 1142–1144
 - stability conditions, 1128–1142
 - Tomita–Takesaki theory, 1119–1122
 - toy-model, 1115–1119
- Kochen–Specker theorem, 359–362, 387, 1405
- Kolmogorov axioms, 644, 1089
- Kosso, P., 1350
- Kraus operator, 570, 756
- Kretschmann, E., 206, 1353
- Kuhn, T., 423

- Lagrange, J. L., 1337, 1345
- Lagrangian approach to mechanics and field theory, 143–146, 154–164, 179, 192
- Lagrangian submanifold, 485
- Landau, L., 417
- Landsman’s theorem, 320
- Lanford’s theorem, 1031–1033
- Laplace’s demon, 1388, 1392
- Laplace, P.-S., 942–943, 1388
- large scale of the universe, 1220
- Last Scattering Surface, *see* LSS
- lattice of propositions, 309–312, 403–405, 641
- laws of creation of the universe, 1218
- laws of nature

- empiricist conception, 1371
- local algebras, 745–750, *see* global observables
- Lebowitz, J., 596
- left action, 56
- left translation, 45
- left-invariant function, 123
- Leggett’s argument against decoherence, 516
- Leibniz, G., 1313
- Lewis, D., 155, 1371
- Lie algebra, 35, 1341
- Lie algebra of a Lie group, 45, 46
- Lie bracket, 37
- Lie derivative, 33
- Lie group, *see* groups and group theory, Lie algebra
- Lie groupoids, 462
- Lie subgroup, 50
- Lie, S., 15
 - and group theory, 1338–1342
- Lie–Poisson bracket, 88, 124, 125
- Lie–Poisson reduction theorem, 19, 77, 82, 88
- limit $\hbar \rightarrow 0$, *see* classical limit
- limit $N \rightarrow \infty$, *see* classical limit
- limit to size of structure, 1205
- Linden, N., 631
- Lipschitz condition, 1383, 1387
- Lo, H.-K., 594, 601
- locality, 308, 351, 381, 391, 396
- localizability of particles, 481
- Lorentz theory, 649, 651
- Lorentz, H., 417
- Lorentz-invariance, 383, 388, 396
- LSS (Last Scattering Surface), 1192, 1202, 1206, 1214, 1230
- Lüder’s rule, 284, 289, *see* collapse postulate
- Mackey’s theorem, 331–334, 447–452, 461–462, *see* system of imprimitivity
- Mackey, G., 447
- macroscopic average, 494
- macroscopic observables, 493–496, 504–506, 512, *see* global observable
- magnetic monopoles, 718
- magnitude-redshift relation, 1214
- Malmquist bias, 1199
- many universes, 1219, 1259–1266
- many-somethings interpretation of quantum mechanics, 370, 646, 1407–1408
- Marsden–Weinstein reduction, 467
- Marsden–Weinstein–Meyer theorem, 127
- matrix mechanics, 277, 417, 427, 428, 641
- matrix representation of quantum states, 282, 298, 364
- Maudlin, T., 183, 184, 191, 209
- Maxwell, J., 944–952, 1300
- Maxwell’s equations, 188–190, 1395, 1416
- Mayers, D., 594, 601, 608
- measurement, 321–325, 327, 328, 334–337, 345, 352, 355–359, 362, 569
 - first kind, 324
 - generalized, 569, 573, 579
 - impulsive, 323, 326, 338, 376
 - non-selective measurement, 394, 575
 - operator, 323, 570–572, 579
 - projective, 570, 571
 - and entropy, 579
 - repeated, 311, 325
 - second kind, 324
 - selective, 575
 - weak, 324–327
- measurement problem, 275, 355–368, 376–381, 431–433, 512, 518, 526, 1163–1168, *see* Schrödinger’s problem
- microcausality, 739, 741–743, 746, 749–753, 766, 771, 772, 807, 839,

- see* kinematic independence
- microlocal analysis, 419, 476, 482
- microscopic quantum observables, 506
- Milne model, 1189
- minimum apparent size, 1198
- Minkowski spacetime, 256–263
- Minkowski, H., 1363
- misconceptions about cosmology, 1214
- Misner, C., 1291
- mixed state, 284, 289, 290, 316, 356, 364, 430–433, 439, 473, 479, 494, 501, 502, 565, 568–571, 574, 576–578, 588–592, 636, 639, 643, 749, *see* ignorance interpretation
 - improper, 289, 366
 - maximally mixed, 749
 - proper, 289, 366, 605
- mixture, *see* mixed state
- modal interpretations of quantum theory, 375, 646, 761–766, 1406–1407
- model selection, 1229
- module, 876
- momentum(or non-relativistic quantum mechanics), 287, 329, 335, 340, 349–352
- momentum function, 114
- momentum map, 6, 106, 464
- momentum map on cotangent bundle, 114
- monic, 874
- monoid, 876
 - absorbing, 902, 913
 - module, 876
 - quotient, 878
- n -Body problem, 175–177, *see* Newtonian n -body problem
- Naimark’s dilation theorem, 453
- naive realism, 359
- natural transformation, 864
 - monoidal, 866
- Navier–Stokes equations, 1383–1384, 1396
- Neumann, J. von, 430, 432
- Newton, I., 1312, 1314
- Newton’s equations of motion, 1385
- Newton–Wigner position operator, 451
- Newtonian n -body problem, 139–146, 1385
- Newtonian gravitation theory, 245, 263–270, 1385
- Nieuwenhuizen, P. van, 1292
- ‘no go’ theorems, 383, 645
- no-signaling theorem, 391
- Noether’s theorems, 4, 6, 11, 108, 161, 253, 1355–1359, 1378, 1411
- noisy channel, 562, 564
- non-baryonic matter, *see* dark matter
- non-contextuality, 360, 387
- non-locality, 308, 381–397
 - in quantum field theory, 750–756
- non-localizability of particles, *see* localizability of particles
- non-singular start of universe, 1211
- Norton, J., 197
- NP**-complete, 615
- nucleosynthesis, 1192, 1199, 1214, 1222, 1232
- null cone observation, 1194
- number count dipole, 1230
- number counts, 1198, 1199, 1214, 1224, 1227
- object
 - conjugate, 869
 - irreducible, 870
 - projective, 875
- observable, 277, 282, 283, 289, 312, 328, 337, 359, *see* POVM, PVM
 - positive-operator-valued measurement, 277, 569–575
- observable universe, 1183, 1198, 1206, 1222, 1230
- observational cosmology theorem, 1222

- observational completeness, 1232
- observed isotropy, 1223
- Ohanian, H., 1343
- one-time pad, 594
- operators introduced, 277f., 401–403
 - at a point in quantum field theory, 766–777
- origin of the light elements, 1193
- origin of the universe, 1212, *see* big bang
- outcome independence, 390, 395

- Painlevé’s conjecture, 1385
- parameter independence, 390–395
- partial observable, 1308, 1317
- particles in quantum field theory, 756–761, *see* Unruh effect
- passage to the quotients, 63, 76
- past distinguishing, *see* causal properties of relativistic spacetimes
- Pauli, W., 443, 555, 1290
- Pearle, P., 379, 381, 645, 1403
- Penrose, R., 1419
- Peres, A., 590, 608, 1293
- perfect (anti-)correlation, 288, 323, 349, 365, 381, 383
- permutation-invariant state, 500
- Perron–Frobenius theorem, 1048, 1050
- phase gate, 616, 632
- phase space quantization, 452, 456
- phenomenon for Bohr, 434, 442
- Phoenix universe, 1213
- physical proposition, 310, 311–312, 316–317
- physics horizon, 1232, 1233
- $\pi/8$ gate, 616, 632
- Piron’s theorem, 312
- Planck scale, 1287, 1289
- Planck, M., 423, 1075–1079, 1370
- pocket universe, 1210, *see* inflationary model, chaotic
- Podolsky, B., 349

- Poincaré’s recurrence theorem, 983–984
- pointer state, 323, 519
- Poisson antimorphism, 104, 105, 108, 115
- Poisson bracket, 84, 307, 318, 319
- Poisson equation, 1379
- Poisson manifold, 6, 14, 24, 79, 84, 317, 458
 - quotient of, 101
- Poisson map, 93
- Poisson reduction theorem, 101
- Poisson space, 506
- Poisson submanifold, 95
- Poisson structure, 91
- polarization spectrum, 1203
- Polyakov, V., 1300
- Popov, V., 1292
- position in non-relativistic quantum mechanics, 330–333, 335–338, 340, 349–352, 375, 382, 451
- POVM (positive operator valued measure), 277–283, 321, 327, 453, 569–575
- pre- and post-selected states, 326, 596
- predictability, 1388, *see* domain of prediction
 - and chaos, 1388–1389
 - and determinism, 1386–1390
 - in general relativistic physics, 1419
 - in quantum mechanics, 1402–1403
 - in special relativistic physics, 1396–1398
- preferred basis, 372
- prequantization, 463
 - line bundle, 464
- presymplectic manifold, 151, 172, 179–180, 189
- primary states, 498
- Primas, H., 444
- Primas–Zurek cut, 518
- prime factorization, 557
- principle vs. constructive theories, 647–652

- probabilities in the Copenhagen interpretation, 440
- probability, 279, 282, 283, 298, 300, 318, 929, 941–943
- problem of time, 133, 136, 137, 196–221, 1316
- projective representation, 301
- projector morphism, 869
- proper map, 64
- protective measurement, *see* measurement, weak
- public-key cryptography, 617
- pure state in quantum mechanics, 284, 285, 289, 293, 296, 301, 302, 312, 314–317, 320–322, 358
- purification, 568, 569, 574
- PVM (projection-valued measure), 278, 279, 281, 331, 571

- quantization, 153, 154, 177, 195, 220–221, 304, 446, 467, 669, *see* geometric quantization *and* system of imprimitivity *and* Weyl quantization
- quantum information and computation, *see* key distribution, bit commitment, information algorithms, 617
 - communication, 556, 590, 592
 - computation, 295, 556, 614
 - computer, 556
 - circuit model, 616
 - cluster state, 616
 - one-way, 616
 - cryptography, 556, 594, 595
 - dense coding, 593
 - information, 565, 576, 577, *see* information
 - operation, 574
 - parallelism, 631
 - qubit, 577
 - teleportation, 556, 590, 591
- quantum cosmology, 1183, 1212, 1294
- quantum ergodic theory, 1092–1094
 - for time, 1102–1106
 - for space, 1106–1114
- quantum gravity, 1211, 1212, 1214, 1287
- quantum logic, 309, 369
- quantum object in Primas' sense, 444
- quantum postulate, 438
- quantum probability theory, 283, 1089
- quantum theory of infinite systems, 492, 496, 738–750, 1094–1102
- quantum world, 434, 526
- quasi-equivalent representations, 498, 732, 745, 843, *see* representation
- quasilocal observables, 496
- quintessence, 1200
- quotient space, 6, 62, 63

- Raggio's theorem, 445
- randomized algorithm, 629
- Raussendorff, R., 616
- Raychaudhuri equation, 1186, 1188, 1190, 1211
- Redhead, M., 191
- redshift-distance relation, 1184, 1196–1198
- reduced density operator, 288, 366, 367, 375, 568
- reduced phase space, 14, 21, 24, 128, 152, 183, 189, 202
- reduced state, *see* reduced density operator
- reduction of dynamics, 102, 126, 129
- Reeh–Schlieder theorem, 742–743
- reference frame, 329, 335, 337, 339, 347
- Regge calculus, 1296
- relationism, 17–18, 1312
- relative configuration space, 18, 21
- relative phase space, 21
- relative states, 370, 568
- relativistic spacetime structure
 - causal structure, *see* causal properties of relativistic space-

- times
- conformal structure, 232–233, 270–271
- projective structure, 232–233
- reliable distance measurement, 1221
- remote steering, 590, 591, 637, *see* biorthogonal decomposition
- renormalization, 683, 705–711, 714
- representation, 57, 409
 - inequivalent representations, 778f.
- right action, 56
- right translation, 45
- right-invariant function, 123
- right-invariant vector field, 47
- rigid body, 15, 61, 80, 90
- Rimini, A., 645
- Roberts, J., 800, 801, 815, 837, 840, 845, 855, 886
- Robertson-Walker cosmological models, 1186–1188, 1207, 1215, 1223–1227, 1231
- robust state, 518
- Rosen, N. 349
- Rosenfeld, L., 1290
- rotation in relativity theory, 249–252
- rotation group, 44, 52, 60, 72, 74, 90, 98
- RSA, 617
- Ruetsche, L., 221
- Ruffini, R., 1343
- Rynasiewicz, R., 206

- Sagnotti, A., 1292
- Saunders, S., 646
- scars, 490
- Schmidt decomposition, *see* biorthogonal decomposition
- Schrödinger, E., 418, 428, 556, 569, 590, 591, 637, 639
- Schrödinger coherent states, 455, 473, 480
- Schrödinger equation, 305, 323, 355, 373, 474, 1399–1400

- Schrödinger picture, 306, *see* Heisenberg picture
- Schrödinger's cat, 355, 364, 418–422, 514, 527, 640, *see* measurement problem
- Schrödinger's problem, 641–644, 651, *see* measurement problem
- Schumacher, B. 585
- Schumacher's quantum source coding theorem (noiseless channel coding theorem), 578, 586, 589
- Schwarz, J., 1298
- selection criterion for representations, 502, 781f.
- self-adjointness, 403, 1400–1402, 1421–1423
- semiclassical analysis, 476
- semiclassical regime, 471
- semidirect product Poisson structure, 461
- Shannon, C., 557, 558
- Shannon entropy, 557, 558, 562, 564, 576, 577
- Shannon's noisy channel coding theorem, 562, 564
- Shannon's source coding theorem (noiseless channel coding theorem), 557, 560, 562, 577
- Shor, P., 557, 617
- Shor's algorithm, 617, 626
- signaling, 373, 391, 392, 582, *see* remote steering
- Simon, D., 617, 623
- Simon's algorithm, 617, 623
- simultaneity in relativity theory, 237, 248, 256–263
- singlet state, 349, 381, 384
- singularity theorems, 1190, 1207
- size of the visible universe, 1220
- small universe, 1203, 1206, 1230, 1231
- Smolin, J., 575
- Smolin, L., 1293
- Soler's theorem, 312

- Sorkin, R., 1296
- source evolution, 1221
- space, public and private in relativity theory, 249–252
- space of states, *see* state space
- spacetime
- classical, 263–270
 - full Newtonian, 1387
 - Leibnizian, 1376
 - Maxwellian, 1377
 - neo-Newtonian, 263–270, 1377
 - relativistic, 230–249, *see* Friedmann-Robertson-Walker model *and* causal properties of
 - Gödel spacetime, 1407–1409
 - Malament-Hogarth spacetimes, 1420–1421
 - Minkowski spacetime, 256–263, 1394, 1396–1398
 - Misner spacetime, 1417
 - negative mass, 1422
 - Reissner-Nordström, 1421
 - relativistic, *see* causal properties of
 - Schwarzschild spacetime, 1422
 - Taub-NUT spacetime, 1417
- spatial coordinates contrasted with body coordinates, 54, 76
- spatial homogeneity of universe, 1185, 1209, 1218, 1223–1226
- spatial isotropy of universe, 1185, 1195, 1223
- special initial conditions, *see* fine tuning
- specific heat of solids, 1080
- spectral action, 1296
- spectral family, 278, *see* PVM
- spectral modal interpretations, 375, *see* modal interpretations of quantum theory
- spectral theorem, 278, *see* PVM
- spectral triple, 1296
- spectrum of the area, 1301
- spherically symmetric universe models, 1223
- spin, 291–298, 320, 336, 375, 450, 462
- spin- $\frac{1}{2}$, 290, 402
- spin-network, 1301
- ‘square root of NOT’ gate, 616
- semiclassical analysis, 472
- STC* (symmetric tensor category), 817, 853, 870, 872, 881–888, 896–900, 905–907, 910–913, 919, 920
- STR (special theory of relativity), 154, 164, 188, 256–263, 397, 1347, 1393–1399
- stable causality, 1370, *see* causal properties of relativistic spacetimes
- standard candles, 1200, 1221
- standard model of particle physics, 683–723, 1287
- standard models of cosmology, 1188, 1223
- start of the universe, *see* big bang
- state space in quantum mechanics, 280–282, 284, 289, 300, 308, 313–322, *see* mixed state, KMS states, state space
- Steane, A., 630, 631, 650
- Stein, H., 258, 259
- Stern–Gerlach device, 291, 385
- Stinespring’s theorem, 456
- stochastic dynamics, 1037–1062
 - approach to equilibrium and objections, 1046–1056
 - ergodic theorem for Markov processes, 1049
 - interventionism, 1037, 1054–1057
 - Markov processes and irreversibility, 1038–1039, 1057–1062
 - Markov processes defined, 1039–1043
- Stone’s theorem, 301, 304
- Stone-von Neumann theorem, 342, 447, 450, 1423
- stress-energy tensor, 1394, 1410, *see*

- energy conditions
- strong energy condition, 1187, 1211
- strong non-disturbance, 353, 354
- structure constants, 37
- structure formation, 1201, 1202, 1207, 1209, 1210, 1214, 1217, 1221, 1228, 1244, 1245
- structure function, 86
- subadditivity inequality, 578
- subcategory, 864
 - monoidal, 865
- submanifold, 40
- submersion, 40
- substantivalism, 18, 1312, 1411, *see* container view of spacetime
- super
 - Hilbert space, 885
 - fiber functor, 886
 - group, 885
- super-horizon structure of the universe, 1206, *see* observable universe
- supernovae, 1194, 1200, 1224, 1227
- superposition, 280, 281, 317, 318, 356–358, 378
- superselection sector, 501–506, 778f.
- supersymmetry, 724, 1299
- surface brightness, 1197
- symmetric algebra, 911
- symmetric sequence, 494
- symmetry, 252–256, 301, 321, 331, 337–339, 867, *see* gauge symmetry, general covariance, Noether's theorems
 - and group theory, 1337–1342
 - and irrelevance, 1360–1361
 - and Leibniz's Principle of Sufficient Reason (PSR), 1333, 1335
 - and objectivity, 1363
 - arguments, 1332
 - broken, 668–669, 693, 839–840, 1334–1336
 - classification of, 1344–1345
 - continuous, 1345
 - discrete, 1345
 - gauge, 1344, 1345
 - geometric, 1344
 - global, 1360, 1361
 - group, 1342
 - Hamiltonian, 12, 58
 - internal/external, 1344
 - Lagrangian, 161
 - local, 1359, 1361
 - meaning of, 1342–1345
 - of laws, 1332
 - of objects, 1332
 - of solutions, 1332, 1336
 - of states, 1332, 1335, 1336
 - quantum field theory, 764–766
 - principles, 1332
 - transformations, active and passive, 1342
 - variational, 161
- symplectic foliation of Poisson manifolds, 91, 95, *see* symplectic leaf
- symplectic form, 8
- symplectic group, 55
- symplectic leaf, 96, 318, 458, 508
- symplectic manifold, 8, 149–151
- symplectic map, 55
- symplectic vector field, 12
- system of covariance, 329
- system of imprimitivity, 331–334, 410, 447–454, 461–462, *see* Mackey's theorem
 - generalized, 454
 - relation to deformation quantization, 461
- 't Hooft, G., 1292, 1296
- Tannaka, T., 883
- TC* (tensor *-category), 870
- temporal orientability, 1409
- tensor
 - category, 865
 - category, braided, 867
 - category, finitely generated, 900

- category, symmetric, 867
 - functor, 866
 - subcategory, 865
 - thermodynamics, 925, 932–941, 994,
 - see* thermodynamic limit
 - thermodynamic limit, 1017–1028, 1159–1160
 - Thomas, R., 221
 - Thomas–Fermi model, 1085
 - time homogeneity, 304
 - time of decoupling, 1191
 - time-dependent systems, 191–196
 - time-energy uncertainty relation, 344
 - Timpson, C., 632
 - Tomita–Takesaki theory, 735–738, *see* KMS states
 - trace, 871
 - transformation group C^* -algebra, 452
 - transformation theory, 1345–1346
 - transition probability, 318–321, 479, 502
 - classical limit, 473
 - transition probability space, 319, 320, 503
 - transport equation, 484
 - tunnelling, 1156–1157, 1402
 - Turing, A., 614
 - Turing computability, 1418, *see* Turing machine
 - Turing machine, 614
 - nondeterministic (probabilistic), 615
 - universal, 556, 614, 1392
 - twist, 872
 - two-slit experiment, 347, *see* double slit experiment
 - typical sequence, 558, 559
- Umdeutung*
- functorial quantization, 471
 - Heisenberg, 427
 - of classical pure states, 474
 - quantization, 446
 - Rieffel, 459
- Schrödinger, 428
 - von Neumann, 430
 - unbounded energies, 1232
 - uncertainty relation, 341–349, 1289
 - algebraic derivation of, 344
 - optical derivation of, 343
 - wavefunctional derivation of, 343
 - uniform thermal histories, 1225
 - uniqueness of the universe, 1216, 1219
 - uniqueness of cosmology, 1232
 - unitarily equivalent representations, 447, 732, 843, *see* representation
 - unitarity of a Hamiltonian flow, 509
 - unitary representation, 301
 - Unruh effect, 1143, 1295, 1303, 1424, *see* Hawking effect
 - Unruh, W., 1295
- Vaidman, L., 597
 - van Fraassen, B., 647
 - van Hove limit, 1135, 1160–1162
 - Veltman, F., 1292
 - Vidal, G., 631
 - von Neumann, J., 430, 432
 - von Neumann algebra, 443, 498, 506, 522, 730, 731, 742, 744–756, 762–766, 782, 802, 806–808, 812, 839, 1101–1105, 1110–1114, 1118–1121, 1126, 1144, 1146, 1148, 1163, *see* C^* -algebra
 - hyperfinite, 569
 - type classification, 733–738
 - von Neumann chain, 432, 515
 - von Neumann entropy, 575, 577, 579, 1091
 - von Neumann’s imprimitivity theorem, 450
 - von Neumann’s infinite tensor product, 504
 - von Plato, J., 962, 967
- Wallace, D., 646
 - wave packet revival, 429, 519

wave-particle duality, 345,
wavefunction, 281, 343, 373, 379–381
weak non-disturbance, 352, 354, *see*
state in quantum mechanics
Weber, T., 645
Weyl algebra, 758
Weyl quantization, 460, 473
Weyl relations, 333, 340–3411
Weyl, H., 233, 234, 1340, 1363, 1369
Wheeler, J., 650, 1291
Wheeler-DeWitt equation, 1293
Wien’s law for blackbody radiation,
423, 1078
Wiesner, S., 556, 593
Wigner function, 479
classical limit, 480
Wigner’s friend, 377, 1164, 1166
Wigner’s theorem, 300
Wigner, E. P., 1362
and his hierarchy, 1359–1360
geometric symmetries, 1344
Wilson, K., 1300
Wittgenstein, L., 514
WKB approximation, 484
Wootters, W., 580, 590

yes–no experiments, 310

Zeilinger, A., 632
Zurek, W., 580


This page intentionally left blank

Handbook of
THE PHILOSOPHY OF SCIENCE

General Editors: DOV M. GABBAY, PAUL THAGARD, AND JOHN WOODS

PHILOSOPHY
of PHYSICS

PART B



Edited by Jeremy Butterfield
and John Earman



Philosophy of Physics

Part B

Handbook of the Philosophy of Science

General Editors

Dov M. Gabbay
Paul Thagard
John Woods

Cover: Photograph of Albert Einstein taken in 1948 at Princeton
by Yousuf Karsh



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON • NEW YORK • OXFORD
PARIS • SAN DIEGO • SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

North-Holland is an imprint of Elsevier



Philosophy of Physics Part B

Edited by

Jeremy Butterfield

All Souls College, University of Oxford, Oxford, UK

and

John Earman

Department of History and Philosophy of Science,
University of Pittsburgh, Pittsburgh, PA, USA



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON • NEW YORK • OXFORD
PARIS • SAN DIEGO • SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

North-Holland is an imprint of Elsevier



North-Holland is an imprint of Elsevier
Radarweg 29, PO Box 211, 1000 AE Amsterdam, The Netherlands
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, UK

First edition 2007

Copyright © 2007 Elsevier B.V. All rights reserved

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means electronic, mechanical, photocopying, recording or otherwise without the prior written permission of the publisher

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone (+44) (0) 1865 843830; fax (+44) (0) 1865 853333; email: permissions@elsevier.com. Alternatively you can submit your request online by visiting the Elsevier web site at <http://elsevier.com/locate/permissions>, and selecting *Obtaining permission to use Elsevier material*

Notice

No responsibility is assumed by the publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made

Library of Congress Cataloging-in-Publication Data

A catalog record for this book is available from the Library of Congress

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN-13 (Set): 978-0-444-51560-5
ISBN-10 (Set): 0-444-51560-7
ISBN-13 (Part A): 978-0-444-53001-1
ISBN-10 (Part A): 0-444-53001-0
ISBN-13 (Part B): 978-0-444-53002-8
ISBN-10 (Part B): 0-444-53002-9

For information on all North-Holland publications visit our website at books.elsevier.com
--

Printed and bound in The Netherlands

07 08 09 10 11 10 9 8 7 6 5 4 3 2 1

GENERAL PREFACE

Dov Gabbay, Paul Thagard, and John Woods

Whenever science operates at the cutting edge of what is known, it invariably runs into philosophical issues about the nature of knowledge and reality. Scientific controversies raise such questions as the relation of theory and experiment, the nature of explanation, and the extent to which science can approximate to the truth. Within particular sciences, special concerns arise about what exists and how it can be known, for example in physics about the nature of space and time, and in psychology about the nature of consciousness. Hence the philosophy of science is an essential part of the scientific investigation of the world.

In recent decades, philosophy of science has become an increasingly central part of philosophy in general. Although there are still philosophers who think that theories of knowledge and reality can be developed by pure reflection, much current philosophical work finds it necessary and valuable to take into account relevant scientific findings. For example, the philosophy of mind is now closely tied to empirical psychology, and political theory often intersects with economics. Thus philosophy of science provides a valuable bridge between philosophical and scientific inquiry.

More and more, the philosophy of science concerns itself not just with general issues about the nature and validity of science, but especially with particular issues that arise in specific sciences. Accordingly, we have organized this Handbook into many volumes reflecting the full range of current research in the philosophy of science. We invited volume editors who are fully involved in the specific sciences, and are delighted that they have solicited contributions by scientifically-informed philosophers and (in a few cases) philosophically-informed scientists. The result is the most comprehensive review ever provided of the philosophy of science.

Here are the volumes in the Handbook:

Philosophy of Science: Focal Issues, edited by Theo Kuipers.

Philosophy of Physics, edited by Jeremy Butterfield and John Earman.

Philosophy of Biology, edited by Mohan Matthen and Christopher Stephens.

Philosophy of Mathematics, edited by Andrew Irvine.

Philosophy of Logic, edited by Dale Jacquette.

Philosophy of Chemistry and Pharmacology, edited by Andrea Woody and Robin Hendry.

Philosophy of Statistics, edited by Prasanta S. Bandyopadhyay and Malcolm Forster.

Philosophy of Information, edited by Pieter Adriaans and Johan van Benthem.

Philosophy of Technological Sciences, edited by Anthonie Meijers.

Philosophy of Complex Systems, edited by Cliff Hooker and John Collier.

Philosophy of Earth Systems Science, edited by Bryson Brown and Kent Peacock.

Philosophy of Psychology and Cognitive Science, edited by Paul Thagard.

Philosophy of Economics, edited by Uskali Mäki.

Philosophy of Linguistics, edited by Martin Stokhof and Jeroen Groenendijk.

Philosophy of Anthropology and Sociology, edited by Stephen Turner and Mark Risjord.

Philosophy of Medicine, edited by Fred Gifford.

Details about the contents and publishing schedule of the volumes can be found at <http://www.johnwoods.ca/HPS/>.

As general editors, we are extremely grateful to the volume editors for arranging such a distinguished array of contributors and for managing their contributions. Production of these volumes has been a huge enterprise, and our warmest thanks go to Jane Spurr and Carol Woods for putting them together. Thanks also to Andy Deelen and Arjen Sevenster at Elsevier for their support and direction.

ACKNOWLEDGEMENTS

It is a pleasure to thank several people. First and foremost, we are grateful to John Woods, Dawn Collins, Jane Spurr and their colleagues at Lethbridge, Alberta, and London, England for their unflagging support and encouragement. We also thank Brian Hepburn for invaluable help with references and the preparation of the text.

This page intentionally left blank

CONTENTS

General Preface	v
Dov Gabbay, Paul Thagard, and John Woods	
Introduction	xiii
Jeremy Butterfield and John Earman	
List of Contributors	xxiii
Part A	
On Symplectic Reduction in Classical Mechanics	1
Jeremy Butterfield	
The Representation of Time and Change in Mechanics	133
Gordon Belot	
Classical Relativity Theory	229
David B. Malament	
Non-Relativistic Quantum Mechanics	275
Michael Dickson	
Between Classical and Quantum	417
N. P. Landsman	
Quantum Information and Computing	555
Jeffrey Bub	
The Conceptual Basis of Quantum Field Theory	661
Gerard 't Hooft	
Algebraic Quantum Field Theory	731
Hans Halvorson (with an Appendix by Michael Müger)	
Index to Part A and Part B	I-1

Part B

Compendium of the Foundations of Classical Statistical Physics	923
Jos Uffink	
Quantum Statistical Physics	1075
G�rard Emch	
Issues in the Philosophy of Cosmology	1183
George F.R. Ellis	
Quantum Gravity	1287
Carlo Rovelli	
Symmetries and Invariances in Classical Physics	1331
Katherine Brading and Elena Castellani	
Aspects of Determinism in Modern Physics	1369
John Earman	
Index to Part A and Part B	I-1

This volume is dedicated to the memories of

Robert Clifton (1964–2002)

and

James Cushing (1937–2002).

This page intentionally left blank

INTRODUCTION

Jeremy Butterfield and John Earman

1 THE PHILOSOPHY OF PHYSICS TODAY

In the last forty years, philosophy of physics has become a large and vigorous branch of philosophy, and so has amply won its place in a series of Handbooks in the philosophy of science. The reasons for its vigour are not far to seek. As we see matters, there are two main reasons; the first relates to the formative years of analytic philosophy of science, and the second to the last forty years.

First, physics had an enormous influence on the early phase of the analytic movement in philosophy. This influence does not just reflect the fact that for the logical positivists and logical empiricists, and for others such as Russell, physics represented a paradigm of empirical knowledge. There are also much more specific influences. Each of the three main pillars of modern physics — thermal physics, quantum theory and relativity — contributed specific ideas and arguments to philosophical debate. Among the more obvious influences are the following.

Thermal physics and the scientific controversy about the existence of atoms bore upon the philosophical debate between realism and instrumentalism; and the rise of statistical mechanics fuelled the philosophy of probability. As to quantum theory, its most pervasive influence in philosophy has undoubtedly been to make philosophers accept that a fundamental physical theory could be indeterministic. But this influence is questionable since, as every philosopher of science knows (or should know!), indeterminism only enters at the most controversial point of quantum theory: viz., the alleged “collapse of the wave packet”. In any case, the obscurity of the interpretation of quantum theory threw not only philosophers, but also the giants of physics, such as Einstein and Bohr, into vigorous debate: and not only about determinism, but also about other philosophical fundamentals, such as the nature of objectivity. Finally, relativity theory, both special and general, revolutionized the philosophy of space and time, in particular by threatening neo-Kantian doctrines about the nature of geometry.

These influences meant that when the analytic movement became dominant in anglophone philosophy, the interpretation of modern physics was established as a prominent theme in its sub-discipline, philosophy of science. Accordingly, as philosophy has grown, so has the philosophy of physics.

But from the 1960s onwards, philosophy of physics has also grown for a reason external to philosophy. Namely, within physics itself there has been considerable

interest in foundational issues, with results that have many suggestive repercussions for philosophy. Again, there have been various developments within physics, and thereby various influences on philosophy. The result, we believe, is that nowadays foundational issues in the fundamental physical theories provide the most interesting and important problems in the philosophy of physics. We have chosen the topics for this volume in accord with this conviction. In the next Subsection, we will articulate some of these foundational issues, and thereby introduce the Chapters of the volume.

2 CURRENT FOUNDATIONAL ISSUES IN PHYSICS

We will first discuss these issues under five headings. The first three correspond to the three pillars of modern physics mentioned in Section 2.1; i.e. thermal physics, quantum theory and relativity theory. The fourth and fifth concern combinations of these pillars; and lead to speculations about the future of physics. These five headings will provide a way of introducing most of this volume's Chapters, albeit not in the order in which they occur. Then, after these five headings, we will introduce the volume's remaining two Chapters.

2.1 *Thermal physics*

Controversies about the foundations of thermal physics, especially the characterization of the approach to equilibrium, have continued unabated since the days of the field's founding fathers, such as Maxwell and Boltzmann. Some aspects of the original controversies can be seen again in modern discussions. But the controversies have also been transformed by the development of several scientific fields; especially the following three, which have grown enormously since the 1960s:

- (i) classical mechanics, and its offspring such as ergodic theory and chaos theory;
- (ii) quantum thermal physics; and
- (iii) cosmology, which nowadays provides a very detailed and so fruitful context for developing and evaluating Boltzmann's bold idea that the ultimate origin of the "arrow of time" is cosmological.

In this volume, the foundations of thermal physics is represented by the Chapters by Uffink and by Emch, who cover classical and quantum aspects, respectively. Among the topics Uffink discusses, two receive special attention: the evolution of Boltzmann's views, and the mathematical framework of stochastic dynamics. Emch adopts the formalism of algebraic quantum statistical mechanics, and reviews many results about that formalism's notion of equilibrium, i.e. KMS states. Two other Chapters also provide a little stage-setting for Uffink and Emch, though without pursuing the relation to thermal physics: viz. the Chapters by Butterfield on classical mechanics, and by Ellis on cosmology.

2.2 *Quantum theory*

Since the 1960s, the physics community has witnessed a revival of the debates about the interpretation of quantum theory that raged among the theory's founding fathers. In the general physics community, the single most influential author has no doubt been John Bell, not only through his non-locality theorem and the many experiments it engendered, but also through his critique of the "Copenhagen orthodoxy" and his sympathy towards the pilot-wave and dynamical collapse heterodoxies. But in more specialist communities, there have been other crucial factors that have animated the debate. Mathematical physicists have developed a deep understanding of the various relations between quantum and classical theories. Since the 1970s, there has been progress in understanding decoherence, so that nowadays, almost all would accept that it plays a crucial role in the emergence of the classical world from quantum theory. And since the 1990s, the burgeoning fields of quantum information and computation have grown out of the interpretative debates, especially the analysis of quantum non-locality.

In this volume, these topics are taken up by Dickson, Landsman and Bub. Dickson surveys the formalism of non-relativistic quantum theory, and some of the main interpretative issues, including empirical content, quantum uncertainty, the measurement problem, and non-locality. For the most part, Landsman reviews from the perspective of mathematical physics the relations between quantum and classical theories. In particular, he discusses various approaches to quantization and the rigorous treatments of the classical limits $\hbar \rightarrow 0$ and $N \rightarrow \infty$. But Landsman also includes discussions of the Copenhagen interpretation and decoherence. Finally, Bub presents some central ideas and results about quantum information and quantum computation. As a backdrop to this, he also briefly reviews classical information and computation; and he ends by proposing some provocative morals about the interpretation of quantum theory.

2.3 *Relativity theory*

The decades since the 1960s have seen spectacular developments, for both theory and experiment, in general relativity and cosmology. But this Renaissance has also been very fruitful as regards foundational and philosophical issues. Mathematical relativists have continued to deepen our understanding of the foundations of general relativity: foundations which, as mentioned in Section 1, were recognized already in the 1920s as crucial for the philosophy of space and time. And the recent transformation of cosmology from a largely speculative enterprise into a genuine science has both brought various philosophical questions closer to scientific resolution, and made other philosophical questions, e.g. about method and explanation in cosmology, much more pressing.

In this volume, these topics are represented by the Chapters by Malament, Belot and Ellis. Malament first expounds classical relativity. Then he discusses three special topics: the definition of simultaneity in special relativity, the geometrization of Newtonian gravity, and the extent to which causal structure determines

spacetime geometry. Belot’s main aim is to give a clear statement of the “problem of time” as it occurs in classical general relativity; and to do that, he first reviews the way time is represented in simpler classical theories, including mechanics. (Belot’s Chapter thereby complements Butterfield’s: both expound aspects of classical Hamiltonian theories, and stress how some of these aspects reappear in quantum theories.) Ellis first reviews the present state of relativistic cosmological theory and its observational basis; and then investigates nine philosophical themes, including the anthropic principle and the possible existence of multiverses.

So much by way of introducing some foundational issues, and this volume’s corresponding Chapters, arising *within* one of the three pillars: thermal physics, quantum theory and relativity. We turn to issues arising from combining the pillars — or rather, parts of them! We have already adumbrated the combination of the first and second: viz., in quantum thermal physics, reviewed here by Emch. It is the combination of the second and third — quantum theory and relativity — which we must now address. We shall do so under two headings, corresponding to the distinction between special and general relativity. The first corresponds, of course, to quantum field theory, which forms such a deep and well-established framework for particle physics. The second corresponds to the quantum theory of gravity — which unfortunately still remains only a hope and a goal.¹

2.4 Quantum field theory

Although there are relativistic quantum mechanical theories of a fixed number of particles, by far the most important framework combining quantum theory and special relativity is quantum field theory. Broadly speaking, the foundational issues raised by quantum field theory differ from quantum theory’s traditional interpretative issues, about measurement and non-locality (cf. *Quantum theory*, §2.2 above). There are two points here.

- (i) Although quantum field theory of course illustrates the latter issues just as much as elementary quantum theory does, it apparently cannot offer a resolution of them. The measurement problem and the puzzles about non-locality arise so directly from the unitarity and tensor-product features of quantum theories, as to be unaffected by the extra mathematical structures

¹Our image of three pillars prompts the question: what about the combination of thermal physics and relativity? When Einstein’s special theory of relativity won acceptance, the rush was on to revise the various branches of classical physics to make them properly relativistic. In the case of thermodynamics, this program produced disputes about the Lorentz transformation properties of the thermodynamic quantities of heat, temperature and entropy that persisted well into the 1970s; (see [Liu, 1994] for an overview of this debate). As for classical general relativity theory, there does not currently exist a statistical mechanics that incorporates the “gravitational entropy of the universe”, and it seems unlikely that there can be such a theory. But for all anyone knows, the ideas of thermal physics may play a crucial role in the hoped-for quantum theory of gravity. There are hints to that effect from, for example, black hole thermodynamics, the Unruh effect, and Hawking radiation. These topics are discussed briefly in Rovelli’s chapter.

and physical ideas supplied by quantum field theory.² And accordingly, it has seemed to most workers to be wisest to pursue the traditional interpretative issues within non-relativistic quantum theory: if you identify a problem in a simple context, but are confident that it is not an artefact of the context's simplicity, it is surely wisest to attack it there. (And as shown in this volume by Dickson's and Landsman's Chapters, that context is by no means "really simple": non-relativistic quantum theory, and its relation to classical theories, provides an abundance of intricate structure to investigate.)

- (ii) On the other hand, there are several foundational issues that are distinctive of quantum field theory. Perhaps the most obvious ones are: the nature of particles (including the topic of localization), the interpretation of renormalization, the interpretation of gauge structure, and the existence of unitarily equivalent representations of the canonical commutation relations.

In this volume, these topics are taken up by 't Hooft and by Halvorson and Müger. First, 't Hooft provides an authoritative survey of quantum field theory, from the perspective of particle physics. Among the main topics he expounds are: the quantization of scalar and spinor fields, Feynman path integrals, the ideas of gauge fields and the Higgs mechanism, renormalization, asymptotic freedom and confinement.

Halvorson and Müger discuss a smaller and distinctively foundational set of issues, using the apparatus of algebraic quantum field theory. (So their use of the algebraic approach complements the uses made by Emch and Landsman.) They discuss the nature of particles and localization, non-locality, the assignment of values to quantities (i.e. the measurement problem) and the definability of quantum fields at individual spacetime points. But they devote most of their effort to the Doplicher-Haag-Roberts theory of superselection. This theory yields deep insights into crucial structures of quantum field theory: in particular, the set of representations, the relation between the field and observable algebras, and gauge groups.

2.5 *Quantum gravity*

Finally, we turn to the combination of quantum theory with general relativity: i.e., the search for a quantum theory of gravity. Here there is of course no established theory, nor even a consensus about the best approach for constructing one. Rather there are various research programmes that often differ in their technical aims, as well as their motivations and conceptual frameworks. In this situation, various

²In some respects relativistic QFT makes the measurement problem worse. In non-relativistic quantum mechanics, the collapse of the state vector is supposed to happen instantaneously; so in the relativistic setting, one would have to develop some appropriate analogue. On the other hand, the modal interpretation of ordinary QM — which arguably provides the best hope for a no-collapse account of quantum measurement — faces formidable obstacles in relativistic quantum field theory; (see [Clifton, 2000] and Halvorson and Müger, this volume, Section 5).

foundational issues about the “ingredient” theories are cast in a new light. For example, might quantum gravity revoke orthodox quantum theory’s unitarity, and thereby *en passant* solve the measurement problem? And does the general covariance (diffeomorphism invariance) of general relativity represent an important clue about the ultimate quantum nature of space and time?

In this volume, these and related questions are taken up by Rovelli. He also presents details about other topics: for example, the subject’s history, the two main current programmes (string theory and loop quantum gravity), and quantum cosmology. Ellis’ Chapter also discusses quantum cosmology. In this way, and indeed by addressing other fundamental questions about the idea of an “ultimate” physical theory, Ellis’s Chapter provides a natural complement to Rovelli’s.

So much by way of introducing Chapters that correspond to our initial three pillars of modern physics, or to combinations of them. We turn to introducing the volume’s remaining two Chapters. Here our intention has been to provide Chapters whose discussions bridge the divisions between physical theories, and even those between our three pillars. In this connection, it seemed to us that of the various possible themes for such a cross-cutting discussion, the two most appropriate ones were determinism and symmetry.³

Accordingly, Earman discusses how determinism fares in a wide class of theories: his examples range from classical mechanics to proposals for quantum gravity. He also addresses the relations between determinism and other issues: in particular, predictability, the nature of spacetime, and symmetry. Symmetry in classical physics is given a wide-ranging survey by Brading and Castellani. Among other topics, they discuss: Curie’s principle, the advent of group theory into physics, canonical transformation theory, general covariance in general relativity, and Noether’s theorems. Various aspects of symmetry and invariance in quantum physics are discussed in the Chapters by Dickson, Emch, Halvorson, and Landsman. But a synoptic overview of this complex topic remains to be written — which we hope will be taken as a challenge by some of our readers.

Let us sum up this introduction to the Chapters that follow, with two comments that are intended to give the prospective reader — perhaps daunted by the many pages ahead! — some courage.

First, it is obvious that by our lights, there is no sharp line between philosophy of physics and physics itself. So it is no surprise that some of the best work in philosophy of physics is being done by physicists (as witnessed by several contributions to this volume). No surprise: but certainly, to be welcomed. Conversely, to the traditionally trained philosopher, work by philosophers of physics is liable to look more like physics than philosophy. But for us, this blurring of disciplinary boundaries is no cause for concern. On the contrary, it represents an opportunity for philosophy to enrich itself. And in the other direction, philosophers can hope

³Other good candidates include the “direction of time”, or irreversibility, and the constitution of matter. But adding chapters on these or other cross-cutting themes would have made the volume altogether too long.

that the foundations, and even philosophy, of physics can be a source of heuristic ideas for physics. Or at least, physicists' interest in foundational questions now offers philosophers of physics the opportunity of fruitful discussion with physicists.

But agreed: this enrichment of philosophy does not come for free. And the need to master technical material which is often difficult can be a barrier to entering the philosophy of physics. In designing this volume, our response to this problem has of course been, not to try to lower the barrier, at the cost of scholarship and of fostering illusory hopes: rather our strategy has been to commission Chapters that cover their chosen topics as expertly and completely as possible. So to the reader, our message is simple: take heart! Once you are over the barrier, new vistas open for the philosophy of science.

3 OUTLOOK: HALFWAY THROUGH THE WOODS

Finally, we would like to set the stage for this volume, by making two connected comments about the present state of fundamental physics. Though it may seem naive or hubristic for philosophers to make such comments, we believe it is worth the risk. For we think that at the present juncture fundamental physics is unusually open to contributions from philosophical reflection; and it will be clear from our comments that together they represent an invitation to the reader to make such contributions! The first comment concerns the amazing successes of present-day physics; the second, the fact that so much remains to be understood.

3.1 Successes

First, we want to celebrate the extraordinary achievements of modern physics; specifically of quantum theory and relativity theory. We propose to do this by emphasising how contingent, indeed surprising, it is that the basic postulates of relativity and quantum theory have proved to be so successful in domains of application far beyond their originally intended ones.

Examples are legion. We pick out two examples, almost at random. Why should the new chronogeometry introduced by Einstein's special relativity in 1905 for electromagnetism, be extendible to mechanics, thermodynamics and other fields of physics? And why should the quantum theory, devised for systems of atomic dimensions (10^{-8} cm) be good both for scales much smaller (cf. the nuclear radius of ca. 10^{-12} cm) and vastly larger (cf. superconductivity and superfluidity, involving scales up to 10^{-1} cm)? Indeed, much of the history of twentieth century physics is the story of the consolidation of the relativity and quantum revolutions: the story of their basic postulates being successfully applied ever more widely.

The point applies equally well when we look beyond terrestrial physics. We have in mind, first, general relativity. It makes a wonderful story: the theory was created principally by one person, motivated by conceptual, in part genuinely philosophical, considerations — yet it has proved experimentally accurate in all

kinds of astronomical situations. They range from weak gravitational fields such as occur in the solar system — here it famously explains the minuscule portion of the precession of the perihelion of Mercury (43" of arc per century) that was unaccounted for by Newtonian theory; to fields 10,000 times stronger in a distant binary pulsar — which in the last twenty years has given us compelling (albeit indirect) evidence for a phenomenon (gravitational radiation) that was predicted by general relativity and long searched for; and to exotic objects such as black holes. But general relativity is not the only case. Quantum theory has also been extraordinarily successful in application to astronomy: the obvious example is the use of nuclear physics to develop a very accurate and detailed theories of nucleosynthesis in the very early universe, and of stellar structure and evolution.

Indeed, there is a more general point here, going beyond the successes of relativity and quantum theory. Namely, we tend to get used to the various unities in nature that science reveals — and thereby to forget how contingent and surprising they are. Of course, this is not just a tendency of our own era. For example, nineteenth century physics confirmed Newton's law of gravitation to apply outside the solar system, and discovered terrestrial elements to exist in the stars (by spectroscopy): discoveries that were briefly surprising, but soon taken for granted, incorporated into the educated person's 'common sense'. Similarly nowadays: the many and varied successes of physics in the last few decades, in modelling very accurately phenomena that are (i) vastly distant in space and time, and-or (ii) very different from our usual laboratory scales (in their characteristic values of such quantities as energy, temperature, or pressure etc.), reveal an amazing unity in nature. General theoretical examples of such unity, examples that span some 200 years, are: the ubiquitous fruitfulness of the field concept; and more specifically, of least action principles. For a modern, specific (and literally spectacular) example, consider the precision and detail of our models of supernovae; as confirmed by the wonderful capacity of modern telescope technology to see and analyse individual supernovae, even in other galaxies.

3.2 Clouds on the horizon

And yet: complacency, let alone triumphalism, is not in order! Current physics is full of unfinished business — that is always true in human enquiry. But more to the point, there are clouds on the horizon that may prove as great a threat to the continued success of twentieth century physics, as were the anomalies confronting classical physics at the end of the nineteenth century.

Of course, people differ about what problems they find worrisome; and among the worrisome ones, about which problems are now ripe for being solved, or at least worth addressing. As philosophers, we are generalists: so we naturally find all the various foundational issues mentioned above worrisome. But being generalists, we will of course duck out of trying to say which are the closest to solution, or which are most likely to repay being addressed! In any case, such judgments are hard to adjudicate, since intellectual temperament, and the happenstance of what one

knows about or is interested in, play a large part in forming them.

But we would like to end by returning to one of Section 2's "clouds": a cloud which clearly invites philosophical reflection, and perhaps contributions. Namely, the problem of quantum gravity; in other words, the fact that general relativity and quantum theory are yet to be reconciled. As mentioned in Section 2.5, Rovelli (this volume) discusses how the contrasting conceptual structures of the "ingredient" theories and the ongoing controversies about interpreting them, make for conflicting basic approaches to quantum gravity.

But we want here to emphasise another reason why we still lack a successful theory, despite great effort and ingenuity. In short, it is that the successes of relativity and quantum theory, celebrated in Comment 3.1 above, conspire to deprive us of the relevant experimental data.

Thus there are general reasons to expect data characteristic of quantum gravity to arise only in a regime of energies so high (correspondingly, distances and times so short) as to be completely inaccessible to us. To put the point in terms of length: the value of the Planck length which we expect to be characteristic of quantum gravity is around 10^{-33} cm. This is truly minuscule: the diameters of an atom, nucleus, proton and quark are, respectively, about 10^{-8} , 10^{-12} , 10^{-13} , and 10^{-16} cm. So the Planck length is as many orders of magnitude from the (upper limit for) the diameter of a quark, as that diameter is from our familiar scale of a centimetre!

We can now see how quantum gravity research is in a sense the victim of the successes of relativity and quantum theory. For those successes suggest that we will not see any "new physics" intimating quantum gravity even at the highest energies accessible to us. The obvious example is quasars: these are typically a few light-days in diameter, and yet have a luminosity 1000 times that of our galaxy (itself 100,000 light-years across, containing a hundred billion stars). They are the most energetic, distant (and hence past!) celestial objects that we observe: they are now believed to be fuelled by massive black holes in their cores. Yet suggestions, current thirty years ago, that their stupendous energies and other properties that we *can* observe, could only be explained by fundamentally new physics, have nowadays given way to acceptance that "conventional physics" describing events *outside* the black hole's event-horizon can do so. (Agreed, we expect the physics deep inside the black hole, in the vicinity of its singularity, to exhibit quantum gravity effects: but if ever a region deserved the name "inaccessible", this is surely one!) So the situation is ironic, and even frustrating: quantum gravity research is a victim of its ingredient theories' success.

In any case, the search for quantum gravity is wide open. In closing, we would like to endorse an analogy of Rovelli's [1997]. He suggests that our present search is like that of the mechanical philosophers such as Galileo and Kepler of the early seventeenth century. Just as they struggled with the clues given by Copernicus and Brahe, *en route* to the synthesis given by Newton, so also we are "halfway through the woods". Of course we should be wary of too grossly simplifying and periodizing the scientific revolution, and *a fortiori* of facile analogies between different

historical situations. Nevertheless, it is striking what a “mixed bag” the doctrines of figures such as Galileo and Kepler turn out to have been, from the perspective of the later synthesis. For all their genius, they appear to us (endowed with the anachronistic benefits of hindsight), to have been “transitional figures”. One cannot help speculating that to some future reader of twentieth century physics, enlightened by some future synthesis of general relativity and quantum theory, our current and recent efforts in quantum gravity will seem strange: worthy and sensible from the authors’ perspective (one hopes), but a hodge-podge of insight and error from the reader’s!

BIBLIOGRAPHY

- [Clifton, 2000] R. Clifton. The modal interpretation of algebraic quantum field theory. *Physics Letters A*, 271, 167-177, 2000.
- [Liu, 1994] C. Liu. Is there a relativistic thermodynamics? A case study in the meaning of special relativity. *Studies in the History and Philosophy of Science*, 25, 983-1004, 1994.
- [Rovelli, 1997] C. Rovelli. Halfway through the woods. In J. Earman and J. Norton (Eds.), *The Cosmos of Science* (pp. 180-223). Pittsburgh: University of Pittsburgh Press and Konstanz: Universitäts Verlag, 1997.

CONTRIBUTORS

Gordon Belot

Department of Philosophy, University of Pittsburgh, 1001 Cathedral of Learning,
Pittsburgh, PA 15260, USA.

gbelot@pitt.edu

Katherine Brading

Department of Philosophy, University of Notre Dame Notre Dame, 100 Malloy
Hall, Indiana 46556, USA.

kbrading@nd.edu

Jeffrey Bub

Department of Philosophy, University of Maryland, College Park, MD 20742, USA.

jbub@umd.edu

Jeremy Butterfield

Trinity College, Cambridge, CB2 1TQ, UK.

jb56@cus.cam.ac.uk

Elena Castellani

Department of Philosophy, University of Florence, via Bolognese 52, 50139, Firenze,
Italy.

elena.castellani@unifi.it

Michael Dickson

Department of Philosophy, University of South Carolina, Columbia, SC 29208,
USA.

dickson@sc.edu

John Earman

Department of History and Philosophy of Science, 1017 Cathedral of Learning,
University of Pittsburgh, Pittsburgh, PA 15260, USA.

jearman+@pitt.edu

George F. R. Ellis

Mathematics Department, University of Cape Town, Rondebosch, Cape Town
8001, South Africa.

ellis@maths.uct.ac.za

G rard Emch

Department of Mathematics, University of Florida, 358 Little Hall, PO Box 118105, Gainesville, FL 32611–8105, USA.

gge@math.ufl.edu

Hans Halvorson

Department of Philosophy, Princeton University, Princeton, NJ 08544, USA.

hhalvors@princeton.edu

Gerard 't Hooft

Institute for Theoretical Physics, Utrecht University, Leuvenlaan 4, 3584 CC Utrecht, The Netherlands, and Spinoza Institute, Postbox 80.195, 3508 TD Utrecht, The Netherlands.

g.thooft@phys.uu.nl

N.P. Landsman

Institute for Mathematics, Astrophysics, and Particle Physics, Radboud Universiteit Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands

landsman@math.ru.nl

David B. Malament

Department of Logic and Philosophy of Science, University of California at Irvine, 3151 Social Science Plaza, Irvine, CA 92697-5100, USA.

dmalamen@uci.edu

Michael M ger

Institute for Mathematics, Astrophysics, and Particle Physics, Radboud Universiteit Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands.

mueger@math.ru.nl

Carlo Rovelli

Centre de Physique Th orique de Luminy, Universit  de la M diterran e, case 907, F-13288 Marseille, France.

rovelli@cpt.univ-mrs.fr

Jos Uffink

Institute for History and Foundations of Science, Utrecht University, PO Box 80 000, 3508 TA Utrecht, The Netherlands.

uffink@phys.uu.nl

COMPENDIUM OF THE FOUNDATIONS OF CLASSICAL STATISTICAL PHYSICS

Jos Uffink

1 INTRODUCTION

It has been said that an advantage of having a mature, formalized version of a theory is that one may forget its preceding history. This saying is certainly true for the purpose of studying the conceptual structure of a physical theory. In a discussion of the foundations of classical mechanics, for example, one need not consider the work of the Parisian scholastics. In the foundations of quantum mechanics, one may start from the von Neumann axioms, and disregard the preceding “old” quantum theory. Statistical physics, however, has not yet developed a set of generally accepted formal axioms, and consequently we have no choice but to dwell on its history.

This is not because attempts to chart the foundations of statistical physics have been absent, or scarce [Ehrenfest and Ehrenfest-Afanassjewa, 1912; ter Haar, 1955; Penrose, 1979; Sklar, 1993; Emch and Liu, 2001]e.g.. Rather, the picture that emerges from such studies is that statistical physics has developed into a number of different schools, each with its own programme and technical apparatus. Unlike quantum theory or relativity, this field lacks a common set of assumptions that is accepted by most of the participants; although there is, of course, overlap. But one common denominator seems to be that nearly all schools claim the founding fathers, Maxwell, Boltzmann and Gibbs as their champions.

Broadly understood, statistical physics may be characterized as a branch of physics intended to describe the thermal behaviour and properties of matter in bulk, i.e. of macroscopic dimensions in relation to its microscopic corpuscular constituents and their dynamics.¹ In this review, we shall only deal with approaches that assume a finite number of microscopic constituents, governed by classical dynamics. (See [Emch, 2006] for a discussion of quantum statistical physics that also addresses infinite systems.)

The above description is deliberately vague; it does not yet specify what thermal behaviour is, and being a characterization in terms of intentions, leaves open by what methods the goals may be achieved. Let us expand a bit. There are two basic

¹The terms “in bulk” and the distinction “micro/macroscopic” should be understood in a relative sense. Thus, statistical physics may apply to a galaxy or nebula, in which the constituent stars are considered as ‘microscopic constituents’.

ingredients in statistical physics. The first is a mechanical model of a macroscopic material system. For example, a gas may be modeled as a system of point particles, or as hard spheres, or as composite objects, etc. Similarly, one may employ lattice models for solids, and so forth. In general, the particulars of the mechanical model, and its dynamics, will depend on the system of interest.

The second ingredient of the theory on which all approaches agree is the introduction of probability and statistical considerations. Sometimes, textbooks explain the need for this ingredient by pointing to the fact that an exact solution of the equations of motion for mechanical models with a large number of degrees of freedom is unfeasible. But this motivation from deficiency surely underestimates the constructive and explanatory role that probability plays in statistical physics. A slightly better motivation, also found in many textbooks, is that even if the dynamical equations could be solved in detail, most of these details would turn out to be irrelevant for the purpose of characterizing the thermal behaviour. There is some truth in this observation, yet it can hardly be satisfactory as it stands. Certainly, not all details about the microdynamics are irrelevant, e.g. in phase transitions, and one naturally wishes for more concrete information about exactly which details are irrelevant and which are not.

One of the foremost foundational problems in statistical physics is thus to specify and to clarify the status of probabilistic assumptions in the theory. As we shall see, this task already leads to a rough distinction between approaches in which probability arises as a notion explicitly defined in mechanical terms (kinetic theory), and approaches in which it is a conceptually independent ingredient (statistical mechanics).

Next, there are ingredients on which much less consensus can be found. Here is a (partial) list:

- Assumptions about the overwhelmingly large number of microscopic constituents (typically of the order of 10^{23} or more).
- An assumption about the erratic nature of the dynamics (e.g. ergodicity).
- The choice of special initial conditions.
- The role of external influences on the system, i.e., assumptions about whether the system is open to the exchange of energy/momentum with its environment, in combination with an assumed sensitivity of the dynamics under such external disturbances.
- Symmetry of macroscopic quantities under permutation of the microscopic constituents.
- Limits in the resolution or experimental accuracy of macroscopic observers.
- Appeal to a time-asymmetric principle of causality.

The role of each of these ingredients in the recipe of statistical physics is controversial. What many “chefs” regard as absolutely essential and indispensable, is argued to be insufficient or superfluous by many others. A major goal in the foundations of statistical physics should therefore lie in an attempt to sort out which subset of the above ideas can be formulated in a precise and coherent manner to obtain a unified and sufficiently general framework for a theory of statistical physics.

Another issue in which the preceding discussion has been vague is what is meant by the thermal behaviour and properties of macroscopic matter. There are two sources on which one may draw in order to delineate this topic. The first is by comparison to other (older) traditions in theoretical physics that have the same goal as statistical physics but do not rely on the two main ingredients above viz. a mechanical model and probabilistic arguments. There are two main examples: thermodynamics and hydrodynamics. The other source, of course, is observation. This provides a rich supply of phenomena, some of which have thus far withstood full theoretical explanation (e.g. turbulence).

Obviously, a measure of success for statistical physics can be found in the question to what extent this approach succeeds in reproducing the results of earlier, non-statistical theories, where they are empirically adequate, and in improving upon them where they are not. Thus, the foundations of statistical physics also provides a testing ground for philosophical ideas about inter-theory relations, like reduction (cf. [Brush, 1977; Sklar, 1993; Batterman, 2002]). However I will not go into this issue. The remainder of this introduction will be devoted to a rough sketch of the four theories mentioned, i.e. thermodynamics, hydrodynamics, kinetic theory and statistical physics.

1.1 Thermodynamics.

Orthodox thermodynamics is an approach associated with the names of Clausius, Kelvin, and Planck. Here, one aims to describe the thermal properties of macroscopic bodies while deliberately avoiding commitment to any hypothesis about the microscopic entities that might constitute the bodies in question. Instead, the approach aims to derive certain general laws, valid for all such bodies, from a restricted set of empirical principles.

In this approach the macroscopic body (or thermodynamic system) is conceived of as a sort of black box, which may interact with its environment by means of work and heat exchange. The most basic empirical principle is that macroscopic bodies when left to themselves, i.e. when isolated from an environment, eventually settle down in an equilibrium state in which no further observable changes occur. Moreover, for simple, homogeneous bodies, this equilibrium state is fully characterized by the values of a small number of macroscopic variables.

Other empirical principles state which types of processes are regarded as impossible. By ingenious arguments one can then derive from these principles the existence of certain quantities (in particular: absolute temperature, energy and

entropy) as ‘state functions’, i.e. functions defined on a space of thermodynamical equilibrium states for all such systems.

While the theory focuses on processes, the description it can afford of such processes is extremely limited. In general, a process will take a system through a sequence of non-equilibrium states, for which the thermodynamic state functions are not defined, and thus cannot be characterized in detail with the tools afforded by the theory. Therefore one limits oneself to the consideration of special types of processes, namely those that begin and end in an equilibrium state. Even more special are those processes that proceed so delicately and slowly that up to an arbitrarily small error one may assume that the system remains in equilibrium throughout the entire process. The latter processes are called *quasistatic*, or sometimes *reversible*.²

Of course, since equilibrium states are by definition assumed to remain in equilibrium if unperturbed, all such processes are triggered by an external intervention such as pushing a piston or removing a partition. For the first type of process, orthodox thermodynamics can only relate the initial and final state. The second type of process can be (approximately) represented as a curve in the equilibrium state space.

The advantage of the approach is its generality. Though developed originally for the study of gases and liquids, by the late nineteenth century, it could be extended to the behaviour of magnets and other systems. Indeed, the independence of hypotheses about its micro-constituents means that the methods of orthodox thermodynamics can also be –and have been– applied to essentially quantum-mechanical systems (like photon gases) or to more exotic objects like black holes (see [Rovelli, 2006]).

With regard to the foundations of statistical physics, two aspects of thermodynamics are of outstanding importance. First, the challenge is to provide a counterpart for the very concept of equilibrium states and to provide a counterpart for the thermodynamic law that all isolated systems not in equilibrium evolve towards an equilibrium state. Secondly, statistical physics should give an account of the Second Law of thermodynamics, i.e. the statement that entropy cannot decrease in an adiabatically isolated system. Obviously, such counterparts will be statistical; i.e. they will hold on average or with high probability, but will not coincide with the unexceptionally general statements of thermodynamics.

1.2 Hydrodynamics

It would be a mistake to believe that the goals of statistical physics are exhausted by reproducing the laws of thermodynamics. There are many other traditions in theoretical physics that provide a much more detailed, yet less general, characterization of thermal behaviour. A concrete example is hydrodynamics or fluid dynamics. In contrast to thermodynamics, hydrodynamics does rely on an assump-

²The reader may be warned, however, that there are many different meanings to the term ‘reversible’ in thermodynamics. See [Uffink, 2001] for a discussion.

tion about microscopic constitution. It models a fluid as a continuous medium or plenum. It is, in modern parlance, a field theory. Moreover it aims to describe the evolution of certain macroscopic quantities in the course of time, i.e. during non-equilibrium processes. As such it is an example of a theory which is much more informative and detailed than thermodynamics, at the price, of course, that its empirical scope is restricted to fluids.

Without going in detail (for a more comprehensive account, see e.g. [Landau and Lifshitz, 1987; de Groot and Mazur, 1961]), hydrodynamics assumes there are three fundamental fields: the mass density $\rho(\vec{x}, t)$, a velocity field $\vec{v}(\vec{x}, t)$ and a temperature field $T(\vec{x}, t)$. There are also three fundamental field equations, which express, in a differential form, the conservation of mass, momentum and energy. Unfortunately, these equations introduce further quantities: the pressure $P(\vec{x}, t)$, the stress tensor $\pi(\vec{x}, t)$, the energy density $u(\vec{x}, t)$, the shear and bulk viscosities η and ζ and thermal conductivity κ , each of which has to be related to the fundamental fields by means of various constitutive relations and equations of state (dependent on the fluid concerned), in order to close the field equations, i.e. to make them susceptible to solution.

The resulting equations are explicitly asymmetric under time reversal. Yet another remarkable feature of hydrodynamics is the fact that the equations can be closed at all. That is: the specification of only a handful of macroscopic quantities is needed to predict the evolution of those quantities. Their behaviour is in other words *autonomous*. This same autonomy also holds for other theories or equations used to describe processes in systems out of equilibrium: for example the theories of diffusion, electrical conduction in metals, the Fourier heat equation etc. In spite of a huge number of microscopic degrees of freedom, the evolution of a few macroscopic quantities generally seems to depend only on the instantaneous values of these macroscopic quantities. Apart from accounting for the asymmetry under time reversal displayed by such theories, statistical physics should also ideally explain this remarkable autonomy of their evolution equations.

1.3 Kinetic theory

I turn to the second group of theories we need to consider: those that do rely on hypotheses or modeling assumptions about the internal microscopic constitution or dynamics of the systems considered. As mentioned, they can be divided into two rough subgroups: kinetic theory and statistical mechanics.

Kinetic theory, also called the kinetic theory of gases, the dynamical theory of gases, the molecular-kinetic theory of heat etc., takes as its main starting point the assumption that systems (gases in particular) consist of molecules. The thermal properties and behaviour are then related in particular to the motion of these molecules.

The earliest modern version of a kinetic theory is Daniel Bernoulli's (1731). Bernoulli's work was not followed by further developments along the same line for almost a century. But it regained new interest in the mid-nineteenth century.

The theory developed into a more general and elaborate framework in the hands of Clausius, Maxwell and Boltzmann. Clausius extended Bernoulli's model by taking into account the collisions between the particles, in order to show that the formidable molecular speeds (in the order of 10^3 m/s) were compatible with relatively slow rates of diffusion. However, he did not develop a systematic treatment of collisions and their effects. It was Maxwell who was the first to realize that collisions would tend to produce particles moving at a variety of speeds, rather than a single common speed, and proceeded to ask how probable the various values of the velocity would be in a state of equilibrium. Maxwell thus introduced the concept of probability and statistical considerations into kinetic theory.

From 1868 onwards, Boltzmann took Maxwell's investigations further. In his famous memoir of 1872 he obtained an equation for the evolution of the distribution function, the Boltzmann equation, and claimed that every non-stationary distribution function for an isolated gas would evolve towards the Maxwellian form, i.e. towards the equilibrium state. However, along the way, Boltzmann had made various assumptions and idealizations, e.g. neglecting the effect of multi-particle collisions, which restrict his derivations' validity to dilute gases, as well as the *Stoßzahlansatz*, developed by Maxwell in 1867, (or 'hypothesis of molecular disorder' as he later called it).

The Boltzmann equation, or variations of this equation, is the physicists' workhorse in gas theory. The hydrodynamical equations can be derived from it, as well as other transport equations. However, it is well known that it is only an approximation, and commonly regarded as a first step in a hierarchy of more detailed equations. But the foremost conceptual problem is its time-asymmetric nature, which highlights the fact that the Boltzmann equation itself could not be derived from mechanics alone. During Boltzmann's lifetime, this led to two famous objections, the reversibility objection (*Umkehrreinwand*) by Loschmidt and the recurrence objection (*Wiederkehrreinwand*) by Zermelo. A third important challenge, only put forward much more recently by [Lanford, 1975], concerns the consistency of the Boltzmann equation with the assumption that the gas system is a mechanical system governed by Hamiltonian dynamics.

1.4 *Statistical mechanics*

There is only a vague borderline between kinetic theory and statistical mechanics. The main distinctive criterion, as drawn by the Ehrenfests (1912) is this. Kinetic theory is what the Ehrenfests call "the older formulation of statistico-mechanical investigations" or "kineto-statistics of the molecule". Here, molecular states, in particular their velocities, are regarded as stochastic variables, and probabilities are attached to such molecular states of motion. These probabilities themselves are determined by the state of the total gas system. They are conceived of either as the relative *number* of molecules with a particular state, or the relative *time* during which a molecule has that state. (Maxwell employed the first option, Boltzmann wavered between the two.) It is important to stress that in both options the

“probabilities” in question are determined by the *mechanical* properties of the gas. Hence there is really no clear separation between mechanical and statistical concepts in this approach.

Gradually, a transition was made to what the Ehrenfests called a “modern formulation of statistico-mechanical investigations” or “kineto-statistics of the gas model”, or what is nowadays known as statistical mechanics. In this latter approach, probabilities are not attached to the state of a molecule but to the state of the entire gas system. Thus, the state of the gas, instead of determining the probability distribution, now itself becomes a stochastic variable.

A merit of this latter approach is that interactions between molecules can be taken into account. Indeed, the approach is not necessarily restricted to gases, but might in principle also be applied to liquids or solids. (This is why the name ‘gas theory’ is abandoned.) The price to be paid however, is that the probabilities themselves become more abstract. Since probabilities are attributed to the mechanical states of the total system, they are no longer determined by such mechanical states. Instead, in statistical mechanics, the probabilities are usually conceived of as being determined by means of an ‘ensemble’, i.e. a fictitious collection of replicas of the system in question. But whatever role one may wish to assign to this construction, the main point is that probability is now an independent concept, no longer reducible to mechanical properties of the system.

It is not easy to pinpoint this transition in the course of the history, except to say that Maxwell’s work in the 1860s definitely belong to the first category, and Gibbs’ book of 1902 to the second. Boltzmann’s own works fall somewhere in the middle ground. His earlier contributions clearly belong to the kinetic theory of gases (although his 1868 paper already applies probability to an entire gas system); while his work after 1877 is usually seen as elements in the theory of statistical mechanics. However, Boltzmann himself never indicated a clear distinction between these two different theories, and any attempt to draw a demarcation at an exact location in his work seems somewhat arbitrary.

From a conceptual point of view, the transition from kinetic gas theory to statistical mechanics poses two main foundational questions. First: on what grounds do we choose a particular ensemble, or the probability distribution characterizing the ensemble? Gibbs did not enter into a systematic discussion of this problem, but only discussed special cases of equilibrium ensembles (i.e. canonical, micro-canonical etc.) for which the probability distribution was stipulated by some special simple form. A second problem is to relate the ensemble-based probabilities to the probabilities obtained in the earlier kinetic approach for a single gas model.

The Ehrenfests [1912] paper was the first to recognize these questions, and to provide a partial answer. Namely: Assuming a certain hypothesis of Boltzmann’s, which they dubbed the *ergodic hypothesis*, they pointed out that for an isolated system the micro-canonical distribution is the unique stationary probability distribution. Hence, if one demands that an ensemble of isolated systems describing thermal equilibrium must be represented by a stationary distribution, the only choice for this purpose is the micro-canonical one. Similarly, they pointed out

that under the ergodic hypothesis, infinite time averages and ensemble averages were identical. This, then, would provide a desired link between the probabilities of the older kinetic gas theory and those of statistical mechanics, at least in equilibrium and in the infinite time limit. Yet the Ehrenfests simultaneously expressed strong doubts about the validity of the ergodic hypothesis. These doubts were soon substantiated when in 1913 Rosenthal and Plancherel proved that the hypothesis was untenable for realistic gas models.

The Ehrenfests' reconstruction of Boltzmann's work thus gave a prominent role to the ergodic hypothesis, suggesting that it played a fundamental and lasting role in his thinking. Although this view indeed produces a more coherent view of his multi-faceted work, it is certainly not historically correct. Boltzmann himself also had grave doubts about this hypothesis, and expressly avoided it whenever he could, in particular in his two great papers of 1872 and 1877b. Since the Ehrenfests, many authors have presented accounts of Boltzmann's work. Particularly important are Klein [1973] and Brush [1976].

Nevertheless, the analysis of the Ehrenfests did thus lead to a somewhat clearly delineated programme for or view about the foundations of statistical physics, in which ergodicity was a crucial feature. The demise of the original ergodic hypothesis did not halt the programme; the hypothesis was replaced by an alternative (weaker) hypothesis, i.e. that the system is '*metrically transitive*' (nowadays, the name 'ergodic' is often used as synonym). What is more, certain mathematical results of Birkhoff and von Neumann (the ergodic theorem) showed that for ergodic systems in this new sense, the desired results could indeed be proven, modulo a few mathematical provisos that at first did not attract much attention.

Thus there arose the ergodic or "standard" view on the foundations of statistical mechanics; (see, e.g. [Khinchin, 1949, p. 44]). On that view, the formalism of statistical mechanics emerges as follows: A concrete system, say a container with gas, is represented as a mechanical system with a very large number of degrees of freedom. All physical quantities are functions of the dynamical variables of the system, or, what amounts to the same thing, are functions on its phase space. However, experiments or observation of such physical quantities do not record the instantaneous values of these physical quantities. Instead, every observation must last a duration which may be extremely short by human standards, but will be extremely long on the microscopic level, i.e. one in which the microstate has experienced many changes, e.g. because of the incessant molecular collisions. Hence, all we can register are *time averages* of the physical quantities over a very long periods of time. These averages are thus empirically meaningful. Unfortunately they are theoretically and analytically obstreperous. Time averages depend on the trajectory and can only be computed by integration of the equations of motion. The expectation value of the phase function over a given ensemble, the *phase average* has the opposite qualities, i.e. it is easy to compute, but not immediately empirically relevant. However, ergodicity ensures that the two averages are equal (almost everywhere). Thus, one can combine the best of both worlds, and identify the theoretically convenient with the empirically meaningful.

While statistical mechanics is clearly a more powerful theory than kinetic theory, it is, like thermodynamics, particularly successful in explaining and modeling gases and other systems in equilibrium. Non-equilibrium statistical mechanics remains a field where extra problems appear.

1.5 *Prospectus*

The structure of this chapter is as follows. In Section 2, I will provide a brief exposition of orthodox thermodynamics, and in subsection 2.2 an even briefer review of some less-than-orthodox approaches to thermodynamics. Section 3 looks at the kinetic theory of gases, focusing in particular on Maxwell's ground-breaking papers of 1860 and 1867, and investigates the meaning and status of Maxwell's probabilistic arguments.

Section 4 is devoted to (a selection of) Boltzmann's works, which, as mentioned above, may be characterized as in between kinetic theory and statistical mechanics. The focus will be on his 1868 paper and his most celebrated papers of 1872 and 1877. Also, the objections from Loschmidt [1877] and Zermelo [1897] are discussed, together with Boltzmann's responses. Our discussion emphasizes the variety of assumptions and methods used by Boltzmann over the years, and the open-endedness of his results: the ergodic hypothesis, the *Stoßzahlansatz*, the combinatorial argument of 1877, and a statistical reading of the *H*-theorem that he advocated in the 1890s.

Next, Section 5 presents an account of Gibbs' [1902] version of statistical mechanics and emphasizes the essential differences between his and Boltzmann's approach. Sections 6 and 7 give an overview of some more recent developments in statistical mechanics. In particular, we review some results in modern ergodic theory, as well as approaches that aim to develop a more systematic account of non-equilibrium theory, such as the BBGKY approach (named after Bogolyubov, Born, Green, Kirkwood and Yvon) and the approach of Lanford. Section 7 extends this discussion for a combination of approaches, here united under the name *stochastic dynamics* that includes those known as 'coarse-graining' and 'interventionism' or 'open systems'. In all cases we shall look at the question whether or how such approaches succeed in a satisfactory treatment of non-equilibrium.

As this prospectus makes clear, the choice of topics is highly selective. There are many important topics and developments in the foundations of statistical physics that I will not touch. I list the most conspicuous of those here together with some references for readers that wish to learn more about them.

- Maxwell's demon and Landauer's principle: [Klein, 1970; Earman and Norton, 1998; 1999; Leff and Rex, 2003; Bennett, 2003; Norton, 2005; Maroney, 2005; Ladyman *et al.*, 2006].
- Boltzmann's work in the 1880s (e.g. on monocyclic systems) [Klein, 1972; 1974; Bierhalter, 1992; Gallavotti, 1999; Uffink, 2005].

- Gibbs' paradox [van Kampen, 1984; Jaynes, 1992; Huggett, 1999; Saunders, 2006].
- Branch systems [Schrödinger, 1950; Reichenbach, 1956; Kroes, 1985; Winsberg, 2004].
- Subjective interpretation of probability in statistical mechanics [Tolman, 1938; Jaynes, 1983; von Plato, 1991; van Lith, 2001a; Balian, 2005].
- Prigogine and the Brussels-Austin school [Obcemea and Brändas, 1983; Batterman, 1991; Karakostas, 1996; Edens, 2001; Bishop, 2004].

2 ORTHODOX THERMODYNAMICS

2.1 *The Clausius-Kelvin-Planck approach*

Thermodynamics is a theory that aims to characterize macroscopic physical bodies in terms of macroscopically observable quantities (typically: temperature, pressure, volume, etc.) and to describe their changes under certain types of interactions (typically exchange of heat or work with an environment).

The classical version of the theory, which evolved around 1850, adopted as a methodological starting point that the fundamental laws of the theory should be independent of any particular hypothesis about the microscopic constitution of the bodies concerned. Rather, they should be based on empirical principles, i.e. boldly generalized statements of experimental facts, not on hypothetical and hence untestable assumptions such as the atomic hypothesis.

The reasons for this methodology were twofold. First, the dominant view on the goal of science was the positivist-empirical philosophy which greatly valued directly testable empirical statements above speculative hypotheses. But the sway of the positivist view was never so complete that physicists avoided speculation altogether. In fact many of the main founders of thermodynamics eagerly indulged in embracing particular hypotheses of their own about the microphysical constitution of matter.

The second reason is more pragmatic. The multitude of microphysical hypotheses and conjectures was already so great in the mid-nineteenth century, and the prospect of deciding between them so dim, that it was a clear advantage to obtain and present results that did not depend on such assumptions. Thus, when Clausius stated in 1857 that he firmly believed in the molecular-kinetic view on the nature of gases, he also mentioned that he had not previously revealed this opinion in order not to mix this conviction with his work on thermodynamics proper [Clausius, 1857, p. 353].³

³The wisdom of this choice becomes clear if we compare his fame to that of Rankine. Rankine actually predated Clausius in finding the entropy function (which he called 'thermodynamic potential'). However, this result was largely ignored due to the fact that it was imbedded in Rankine's rather complicated theory of atomic vortices.

Proceeding somewhat ahistorically,⁴ one might say that the first central concept in thermodynamics is that of *equilibrium*. It is taken as a fact of experience that macroscopic bodies in a finite volume, when left to themselves, i.e. isolated from an environment eventually settle down in a stationary state in which no further observable changes occur (the ‘Minus First Law’, cf. page 939). This stationary state is called a (thermal) *equilibrium state*. Moreover, for simple, homogeneous bodies, this state is fully characterized by the values of a small number of macroscopic variables. In particular, for fluids (i.e. gases or liquids), two independent variables suffice to determine the equilibrium state.

For fluids, the three variables pressure p , temperature θ and volume V , are thus related by a so-called equation of state, where, following Euler, it has become customary to express pressure as a function of the two remaining variables:

$$(1) \quad p = p(\theta, V)$$

The form of this function differs for different fluids; for n moles of an ideal gas it is given by:

$$(2) \quad p(\theta, V) = nR\theta/V$$

where R is the gas constant and θ is measured on the gas thermometer scale.

The content of thermodynamics developed out of three ingredients. The first is the science of calorimetry, which was already developed to theoretical perfection in the eighteenth century, in particular by Joseph Black [Fox, 1971; Truesdell, 1980; Chang, 2003; 2004]. It involved the study of the thermal changes in a body under the addition of or withdrawal of heat to the system. Of course, the (silent) presupposition here is that this process of heat exchange proceeds so delicately and slowly that the system may always be regarded as remaining in equilibrium. In modern terms, it proceeds ‘quasi-statically’. Thus, the equation of state remains valid during the process.

The tools of calorimetry are those of differential calculus. For an infinitesimal increment dQ of heat added to a fluid, one puts

$$(3) \quad dQ = c_V d\theta + \Lambda_\theta dV,$$

where c_V is called the heat capacity at constant volume and Λ_θ the latent heat at constant temperature. Both c_V and Λ_θ are assumed to be functions of θ and V . The notation d is used to indicate that the heat increment dQ is not necessarily an exact differential, i.e. Q is not assumed to be a function of state.

The total heat Q added to a fluid during a process can thus be expressed as a line integral along a path \mathcal{P} in the (θ, V) plane

$$(4) \quad Q(\mathcal{P}) = \int_{\mathcal{P}} dQ = \int_{\mathcal{P}} (c_V d\theta + \Lambda_\theta dV)$$

⁴I refer to [Uffink, 2001] for more details.

A treatment similar to the above can be given for the quasistatic heat exchange of more general thermal bodies than fluids. Indeed, calorimetry was sufficiently general to describe phase transitions, say from water to ice, by assuming a discontinuity in Λ_θ .

All this is independent of the question whether heat itself is a substance or not. Indeed, Black himself wished to remain neutral on this issue. Even so, much of the terminology of calorimetry somehow invites the supposition that heat is a substance, usually called caloric, and many eighteenth and early nineteenth century authors adopted this view [Fox, 1971]. In such a view it makes sense to speak of the amount of heat contained in a body, and this would entail that dQ must be an exact differential (or in other words: $Q(\mathcal{P})$ must be the same for all paths \mathcal{P} with the same initial and final points). But this turned out to be empirically false, when the effects of the performance of work were taken into account.

Investigations in the 1840s (by Joule and Mayer among others) led to the conviction that heat and work are “equivalent”; or somewhat more precisely, that in every cyclic process \mathcal{C} , the amount of heat $Q(\mathcal{C})$ absorbed by the system is proportional to the amount of work performed by the system. Or, taking $W(\mathcal{C})$ as positive when performed *on* the system :

$$(5) \quad JQ(\mathcal{C}) + W(\mathcal{C}) = 0$$

where $J \approx 4.2\text{Nm/Cal}$ is Joule’s constant, which modern convention takes equal to 1. This is the so-called *First Law of thermodynamics*.

For quasistatic processes this can again be expressed as a line integral in a state space Ω_{eq} of thermodynamic equilibrium states

$$(6) \quad \oint_{\mathcal{C}} (dQ + dW) = 0$$

where

$$(7) \quad dW = -pdV.$$

Assuming the validity of (6) for all cyclic paths in the equilibrium state space implies the existence of a function U on Ω_{eq} such that

$$(8) \quad dU = dQ + dW.$$

The third ingredient of thermodynamics evolved from the study of the relations between heat and work, in particular the efficiency of heat engines. In 1824, Carnot obtained the following theorem.

CARNOT’S THEOREM: Consider any system that performs a cyclic process \mathcal{C} during which (a) an amount of heat $Q_+(\mathcal{C})$ is absorbed from a heat reservoir at temperature θ_+ , (b) an amount of heat $Q_-(\mathcal{C})$ is given off to a reservoir at a temperature θ_- , with $\theta_- < \theta_+$, (c) there is no heat exchange at other stages of the cycles, and (d) some work

$W(\mathcal{C})$ is done on a third body. Let $\eta(\mathcal{C}) := \frac{W(\mathcal{C})}{Q_+(\mathcal{C})}$ be the efficiency of the cycle. Then:

- (1) All quasistatic cycles have the same efficiency. This efficiency is a universal function of the two temperatures, i.e.,
- (9) $\eta(\mathcal{C}) = \eta(\theta_+, \theta_-)$.
- (2) All other cycles have a efficiency which is less or equal to that of the quasi-static cycle.

Carnot arrived at this result by assuming that heat was a conserved substance (and thus: $Q_+(\mathcal{C}) = Q_-(\mathcal{C})$ for all \mathcal{C}), as well as a principle that excluded the construction of a perpetuum mobile (of the first kind).

In actual fact, Carnot did not use the quasistatic/non-quasistatic dichotomy to characterize the two parts of his theorem.⁵

In fact, he used two different characterizations of the cycles that would produce maximum efficiency. (a): In his proof that Carnot cycles belong to class (1), the crucial assumption is that they “might have been performed in an inverse direction and order” [Carnot, 1824, p. 11]. But a little later (p. 13), he proposed a necessary and sufficient condition for a cycle to produce maximum efficiency, namely (b): In all stages which involve heat exchange, only bodies of equal temperature are put in thermal contact, or rather: their temperatures differ by a vanishingly small amount.

Carnot’s theorem is remarkable since it did not need any assumption about the nature of the thermal system on which the cycle was carried out. Thus, when his work first became known to the physics community (Thomson, later known as Lord Kelvin, 1848) it was recognized as an important clue towards a general theory dealing with both heat and work exchange, for which Kelvin coined the name ‘thermodynamics’. Indeed, Kelvin already showed in his first paper (1848) on the subject that Carnot’s universal function η could be used to devise an absolute scale for temperature, i.e. one that did not depend on properties of a particular substance.

Unfortunately, around the very same period it became clear that Carnot’s assumption of the conservation of heat violated the First Law. In a series of papers Clausius and Kelvin re-established Carnot’s theorem on a different footing (i.e. on the first law (5) or, in this case $Q_+(\mathcal{C}) = Q_-(\mathcal{C}) + W(\mathcal{C})$, and a principle that excluded perpetual motion of the second kind) and transformed his results into general propositions that characterize general thermodynamical systems and their changes under the influence of heat and work. For the most part, these investigations were concerned with the first part of Carnot’s theorem only. They led to what is nowadays called the first part of the Second Law; as follows.

First, Kelvin reformulated his 1848 absolute temperature scale into a new one, $T(\theta)$, in which the universal efficiency could be expressed explicitly as:

⁵Indeed, [Truesdell, 1980] argues that this characterization of his theorem is incorrect. See [Uffink, 2001] for further discussions.

$$(10) \quad \eta(T_+, T_-) = 1 - \frac{T_-}{T_+},$$

where $T_i = T(\theta_i)$. Since the efficiency η is also expressed by $W/Q_+ = 1 - (Q_-/Q_+)$, this is equivalent to

$$(11) \quad \frac{Q_-}{T_-} = \frac{Q_+}{T_+}.$$

Next, changing the sign convention to one in which Q is positive if absorbed and negative if given off by the system, and generalizing for cycles in which an arbitrary number of heat reservoirs are involved, one gets:

$$(12) \quad \sum_i \frac{Q_i}{T_i} = 0.$$

In the case where the system is taken through a quasistatic cycle in which the heat reservoirs have a continuously varying temperature during this cycle, this generalizes to

$$(13) \quad \oint_C \frac{dQ}{T} = 0.$$

Here, T still refers to the temperature of the heat reservoirs with which the system interacts, not to its own temperature. Yet Carnot's necessary and sufficient criterion of reversibility itself requires that during all stages of the process that involve heat exchange, the temperatures of the heat reservoir and system should be equal. Hence, in this case one may equate T with the temperature of the system itself.

The virtue of this result is that the integral (13) can now be entirely expressed in terms of quantities of the system. By a well-known theorem, applied by Clausius in 1865, it follows that there exists a function, called entropy S , defined on the equilibrium states of the system such that

$$(14) \quad S(s_1) - S(s_2) = \int_{s_1}^{s_2} \frac{dQ}{T}$$

or, as it more usually known:

$$(15) \quad \frac{dQ}{T} = dS.$$

This result is frequently expressed as follows: dQ has an *integrating divisor* (namely T): division by T turns the inexact (incomplete, non-integrable) differential dQ into an exact (complete, integrable) differential. For one mole of ideal gas (i.e. a fluid for which c_V is constant, Λ_θ vanishes and the ideal gas law (2) applies), one finds, for example:

$$(16) \quad S(T, V) = c_V \ln T + R \ln V + \text{const.}$$

The existence of this entropy function also allows for a convenient reformulation of the First Law for quasistatic processes (8) as

$$(17) \quad dU = TdS - pdV,$$

now too expressed in terms of properties of the system of interest.

However important this first part of the Second Law is by itself, it never led to much dispute or controversy. By contrast, the extension of the above results to cover the second part of Carnot's theorem gave rise to considerably more thought, and depends also intimately on what is understood by '(ir)reversible processes'.

The second part of Carnot's theorem was at first treated in a much more step-motherly fashion. Clausius' [1854] only devoted a single paragraph to it, obtaining the result that for "irreversible" cycles

$$(18) \quad \oint \frac{dQ}{T} \leq 0.$$

But this result is much less easy to apply, since the temperature T here refers to that of the heat reservoir with which the system is in contact, not (necessarily) that of the system itself.

Clausius put the irreversible processes in a more prominent role in his 1865 paper. If an irreversible cyclic process consists of a general, i.e. possibly non-quasistatic stage, from s_i to s_f , and a quasistatic stage, from s_f back to s_i , one may write (18) as

$$(19) \quad \int_{s_i}^{s_f} \frac{dQ}{T} \Big|_{\text{non-qs}} + \int_{s_f}^{s_i} \frac{dQ}{T} \Big|_{\text{qs}} \leq 0.$$

Applying (14) to the second term in the left hand side, one obtains

$$(20) \quad \int_{s_i}^{s_f} \frac{dQ}{T} \Big|_{\text{non-qs}} \leq S(s_f) - S(s_i)$$

If we assume moreover that the generally non-quasistatic process is adiabatic, i.e. $dQ = 0$, the result is

$$(21) \quad S(s_i) \leq S(s_f).$$

In other words, in any adiabatic process the entropy of the final state cannot be less than that of the initial state.

Remarks: 1. The notation \oint for cyclic integrals, and d for inexact differentials is modern. Clausius, and Boltzmann after him, would simply write $\int \frac{dQ}{T}$ for the left-hand side of (13) and (18).

2. An important point to note is that Clausius' formulation of the Second Law, strictly speaking, does not require a general monotonic increase of entropy for any adiabatically isolated system in the course of time. Indeed, in orthodox

thermodynamics, entropy is defined only for equilibrium states. Therefore it is meaningless within this theory to ask how the entropy of a system changes during a non-quasistatic process. All one can say in general is that when a system starts out in an equilibrium state, and ends, after an adiabatic process, again in an equilibrium state, the entropy of the latter state is not less than that of the former.

Still, the Second Law has often been understood as demanding continuous monotonic increase of entropy in the course of time, and often expressed, for adiabatically isolated systems, in a more stringent form

$$(22) \quad \frac{dS}{dt} \geq 0.$$

There is, however, no basis for this demand in orthodox thermodynamics.

3. Another common misunderstanding of the Second Law is that it would only require the non-decrease of entropy for processes in *isolated* systems. It should be noted that this is only part of the result Clausius derived: the Second Law holds more generally for *adiabatic* processes, i.e., processes during which the system remains adiabatically insulated. In other words, the system may be subject to arbitrary interactions with the environment, except those that involve heat exchange. (For example: stirring a liquid in a thermos flask, as in Joule's 'paddle wheel' experiment.)

4. Another point to be noted is that Clausius' result that the entropy in an adiabatically isolated system can never decrease is derived from the *assumption* that one can find a quasistatic process that connects the final to the initial state, in order to complete a cycle. Indeed, if such a process did not exist, the entropy difference of these two states would not be defined. The existence of such quasistatic processes is not problematic in many intended applications (e.g. if s_f and s_i are equilibrium states of a fluid); but it may be far from obvious in more general settings (for instance if one considers processes far from equilibrium in a complex system, such as a living cell). This warning that the increase of entropy is thus conditional on the existence of quasistatic transitions has been pointed out already by [Kirchhoff, 1894, p. 69].

5. Apart from the well-known First and Second Laws of thermodynamics, later authors have identified some more basic assumptions or empirical principles in the theory that are often assumed silently in traditional presentations — or sometimes explicitly but unnamed — which may claim a similar fundamental status.

The most familiar of these is the so-called *Zeroth Law*, a term coined by [Fowler and Guggenheim, 1939]. To introduce this, consider the *relation of thermal equilibrium*. This is the relationship holding between the equilibrium states of two systems, whenever it is the case that the composite system, consisting of these two systems, would be found in an equilibrium state if the two systems are placed in direct thermal contact — i.e., an interaction by which they are only allowed to exchange heat. The zeroth law is now that the assumption that this is a *transitive* relationship, i.e. if it holds for the states of two bodies A and B , and also for the

states of bodies B and C , it likewise holds for bodies A and C .⁶

2.2 Less orthodox versions of thermodynamics

Even within the framework of orthodox thermodynamics, there are approaches that differ from the Clausius-Kelvin-Planck approach. The foremost of those is undoubtedly the approach developed by Gibbs in 1873–1878 [Gibbs, 1906]. Gibbs' approach differs much in spirit from his European colleagues. No effort is devoted to relate the existence or uniqueness of the thermodynamic state variables U , T or S to empirical principles. Their existence is simply assumed. Also, Gibbs focused on the description of equilibrium states, rather than processes.

Previous authors usually regarded the choice of variables in order to represent a thermodynamic quantity as a matter of convention, like the choice of a coordinate system on the thermodynamic (equilibrium) state space. For a fluid, one could equally well choose the variables (p, V) , (V, T) , etc., as long as they are independent and characterize a unique thermodynamic equilibrium state.⁷ Hence one could equally well express the quantities U , S , etc. in terms of any such set of variables. However, Gibbs had the deep insight that some choices are 'better' than others, in the sense that if, e.g., the entropy is presented as a function of energy and volume, $S(U, V)$, (or energy as a function of entropy and volume, $U(S, V)$) all other thermodynamic quantities could be determined from it, while this is generally not true for other choices. For example, if one knows only that for one mole of gas $S(T, V)$ is given by (2), one cannot deduce the equations of state $p = RT/V$ and $U = c_V T$. In contrast, if the function $S(U, V) = c_V \ln U + R \ln V + \text{const.}$ is given, one obtains these equations from its partial derivatives: $\frac{p}{T} = \left(\frac{\partial S}{\partial V}\right)_U$ and $\frac{1}{T} = \left(\frac{\partial S}{\partial U}\right)_V$.

For this reason, Gibbs called

$$(23) \quad dU = TdS - pdV \quad \text{or} \quad dS = \frac{1}{T}dU + \frac{p}{T}dV$$

the *fundamental equation*.⁸ Of course this does not mean that other choices of variables are inferior. Instead, one can find equivalent fundamental equations for

⁶Actually, transitivity alone is not enough. The assumption actually needed is that thermal equilibrium is an *equivalence* relation, i.e., it is transitive, reflective and symmetric (cf. [Boyling, 1972, p. 45]). The idea of elevating this to a fundamental 'Law', is that this assumption, which underlies the concept of temperature, can only be motivated on empirical grounds.

Another such assumption, again often stated but rarely named, is that any system contained in a finite volume, if left to itself, tends to evolve towards an equilibrium state. This has also sometimes been called a 'zereth law' (cf. [Uhlenbeck and Ford, 1963, p.5; Lebowitz, 1994, p. 135]) in unfortunate competition with Fowler & Guggenheim's nomenclature. The name *Minus First Law* has therefore been proposed by [Brown and Uffink, 2001]. Note that this assumption already introduces an explicitly time-asymmetric element, which is deeper than — and does not follow from — the Second Law. However, most nineteenth (and many twentieth) century authors did not appreciate this distinction, and as we shall see below, this Minus First Law is often subsumed under the Second Law.

⁷The latter condition may well fail: A fluid like water can exist at different equilibrium states with the same p, V , but different T [Thomsen and Hartka, 1962]

⁸Note how Gibbs' outlook differs here from the Clausius-Kelvin-Planck view: These authors

such pairs of variables too, in terms of the Legendre transforms of U . (Namely: the Helmholtz free energy $F = U - TS$ for the pair (T, V) ; the enthalpy $U + pV$ for (p, S) , and the Gibbs free energy $U + pV - TS$ for (p, T) .) Further, Gibbs extended these considerations from homogeneous fluids to heterogeneous bodies, consisting of several chemical components and physical phases.

Another major novelty is that Gibbs proposed a variational principle to distinguish stable from neutral and unstable equilibria. (Roughly, this principle entails that for stable equilibrium the function $S(U, V)$ should be concave.) This criterium serves to be of great value in characterizing phase transitions in thermodynamic systems, e.g. the Van der Waals gas (Maxwell used it to obtain his famous “Maxwell construction” or equal area rule [Klein, 1978]). Gibbs work also proved important in the development of chemical thermodynamics, and physical chemistry.

Another group of approaches in orthodox thermodynamics is concerned particularly with creating a more rigorous formal framework for the theory. This is often called *axiomatic thermodynamics*. Of course, choosing to pursue a physical theory in an axiomatic framework does not by itself imply any preference for a choice in its physical assumptions or philosophical outlook. Yet the Clausius-Kelvin-Planck approach relies on empirical principles and intuitive concepts that may seem clear enough in their relation to experience — but are often surprisingly hard to define. Hence, axiomatic approaches tend to replace these empirical principles by statements that are conceptually more precise, but also more abstract, and thus arguably further removed from experience. The first example of this work is Carathéodory 1909. Later axiomatic approaches were pursued, among others, by [Giles, 1964]; [Boyling, 1972]; Jauch [1972; 1975], and by [Lieb and Yngvason, 1999]. All these approaches differ in their choice of primitive concepts, in the formulation of their axioms, and hence also in the results obtained and goals achieved. However, in a rough sense, one might say they all focus particularly on demonstrating under what conditions one might guarantee the mathematical existence and uniqueness of entropy and other state functions within an appropriate structure.

Since the 1940s a great deal of work has been done on what is known as “*non-equilibrium thermodynamics*” or “thermodynamics of irreversible processes” (see e.g. [de Groot, 1951; Prigogine, 1955; de Groot and Mazur, 1961, Yourgrau, 1966; Truesdell, 1969; Müller, 2003]). This type of work aims to extend orthodox thermodynamics into the direction of a description of systems in non-equilibrium states. Typically, one postulates that thermodynamic quantities are represented as continuously variable fields in space and time, with equilibrium conditions holding approximately within each infinitesimal region within the thermodynamic system.

would look upon (23) as a statement of the first law of thermodynamics, interpreting the differentials as infinitesimal increments during a quasistatic *process*, cf. (17). For Gibbs, on the other hand, (23) does not represent a process but a differential equation on the thermodynamic state space whose solution $U(S, V)$ or $S(V, U)$ contains *all* information about the equilibrium properties of the system, including the equations of state, the specific and latent heat, the compressibility, etc. — much more than just First Law.

Again, it may be noted that workers in the field seem to be divided into different schools (using names such as “extended thermodynamics”, “generalized thermodynamics”, “rational thermodynamics”, etc.) that do not at all agree with each other (see [Hutter and Wang, 2003]).

This type of work has produced many successful applications. But it seems fair to say that until now almost all attention has gone to towards practical application. For example, questions of the type that axiomatic thermodynamics attempts to answer, (e.g.: Under what conditions can we show the existence and uniqueness of the non-equilibrium quantities used in the formalism?) are largely unanswered, and indeed have given rise to some scepticism (cf. [Meixner, 1969; Meixner, 1970]). Another inherent restriction of this theory is that by relying on the assumption that non-equilibrium states can, at least in an infinitesimal local region, be well approximated by an equilibrium state, the approach is incapable of encompassing systems that are very far from equilibrium, such as in turbulence or living cells.)

The final type of approach that ought to be mentioned is that of *statistical thermodynamics*. The basic idea here is that while still refraining from introducing hypotheses about the microscopic constituents of thermodynamic systems, one rejects a key assumption of orthodox thermodynamics, namely, that a state of equilibrium is one in which all quantities attain constant values, in order to accommodate fluctuation phenomena such as Brownian motion or thermal noise. Thus the idea becomes to represent at least some of the thermodynamic quantities as random quantities, that in the course of time attain various values with various probabilities. Work in this direction has been done by Szilard [1925], Mandelbrot [1956; 1962; 1964], and Tisza and Quay [1963].

Of course the crucial question is then how to choose the appropriate probability distributions. One approach, elaborated in particular by [Tisza, 1966], taking its inspiration from Einstein [1910], relies on an inversion of Boltzmann’s principle: whereas Boltzmann argued (within statistical mechanics) that the thermodynamic notion of entropy could be identified with the logarithm of a probability; Einstein argued that in thermodynamics, where the concept of entropy is already given, one may define the relative probability of two equilibrium states by the exponent of their entropy difference. Other approaches have borrowed more sophisticated results from mathematical statistics. For example, Mandelbrot used the Pitman-Koopman-Darmois theorem, which states that sufficient estimators exist only for the “exponential family” of probability distributions to derive the canonical probability distribution from the postulate that energy be a sufficient estimator of the system’s temperature (see also [Uffink and van Lith, 1999]).

3 KINETIC THEORY FROM BERNOULLI TO MAXWELL

3.1 *Probability in the mid-nineteenth century*

Probability theory has a history dating back at least two centuries before the appearance of statistical physics. Usually, one places the birth of this theory in

the correspondence of Pascal and Fermat around 1650. It was refined into a mature mathematical discipline in the work of Jacob Bernoulli [1713], Abraham de Moivre [1738] and Pierre-Simon de Laplace [1813] (cf. [Hacking, 1975]).

In this tradition, often called ‘classical probability’, the notion of probability is conceived of as a measure of the degree of certainty of our beliefs. Two points are important to note here. First, in this particular view, probability resides in the mind. There is nothing like uncertainty or chance in Nature. In fact, all authors in the classical tradition emphasize their adherence to strict determinism, either by appeal to divine omniscience (Bernoulli, de Moivre) or by appeal to the laws of mechanics and the initial conditions (Laplace). A probability hence represents a judgment about some state of affairs, and not an intrinsic property of this state of affairs. Hence, the classical authors never conceived that probability has any role to play in a description of nature or physical processes as such.⁹ Secondly, although Bernoulli himself used the term “subjective” to emphasize the fact that probability refers to us, and the knowledge we possess, the classical interpretation does not go so far as modern adherents to a subjective interpretation of probability who conceive of probability as the degrees of belief of an arbitrary (although coherent) person, who may base his beliefs on personal whims, prejudice and private opinion.

This classical conception of probability would, of course, remain a view without any bite, if it were not accompanied by some rule for assigning values to probabilities in specific cases. The only such available rule is the so-called ‘principle of insufficient reason’: whenever we have no reason to believe that one case rather than another is realized, we should assign them equal probabilities (cf. [Uffink, 1995]). A closely related version is the rule that two or more variables should be independent whenever we have no reason to believe that they influence each other.

While the classical view was the dominant, indeed the only existent, view on probability for the whole period from 1650 to 1813, it began to erode around 1830. There are several reasons for this, but perhaps the most important is, paradoxically, the huge success with which the theory was being applied to the most varied subjects. In the wake of Laplace’s influential *Essai philosophique sûr les Probabilités*, scientists found applications of probability theory in jurisdiction, demography, social science, hereditary research, etc. In fact, one may say: almost everywhere except physics (cf. [Hacking, 1990]). The striking regularity found in the frequencies of mass phenomena, and observations that (say) the number of raindrops per second on a tile follows the same pattern as the number of soldiers in the Prussian army killed each year by a kick from their horse, led to the alternative view that probability was not so much a representation of subjective (un)certainly, but rather the expression of a particular regularity in nature (Poisson, Quetelet). From these days onward we find mention of the idea of *laws of probability*, i.e. the idea that theorems of probability theory reflect lawlike behaviour to which

⁹Daniel Bernoulli might serve as an example. He was well acquainted with the work on probability of his uncle Jacob and, indeed, himself one of the foremost probabilists of the eighteenth century. Yet, in his work on kinetic gas theory (to be discussed in section 3.2), he did not find any occasion to draw a connection between these two fields of his own expertise.

Nature adheres. In this alternative, frequentist view of probability, there is no obvious place for the principle of insufficient reason. Instead, the obvious way to determine the values of probabilities is to collect empirical data on the frequencies on occurrences of events. However, a well-articulated alternative to the classical concept of probability did not emerge before the end of the century, and (arguably) not before 1919 — and then within in a few years there were no less than three alternatives: a logical interpretation by Keynes, a frequentist interpretation by von Mises and a subjective interpretation by Ramsey and De Finetti. See [Fine, 1973], [Galavotti, 2004] or [Emch, 2005] for a more detailed exposition.

Summing up roughly, one may say that around 1850 the field of probability was in a state of flux and confusion. Two competing viewpoints, the classical and the frequency interpretation, were available, and often mixed together in a confusing hodgepodge. The result was well-characterized in a famous remark of [Poincaré, 1896] that all mathematicians seem to believe that the laws of probability refer to statements learned from experience, while all natural scientists seem to think they are theorems derived by pure mathematics.

The work of Maxwell and Boltzmann in the 1860s emerged just in the middle of this confusing era. It is only natural that their work should reflect the ambiguity that the probability concept had acquired in the first half of the nineteenth century. Nevertheless, it seems that they mainly thought of probability in terms of frequencies, as an objective quantity, which characterizes a many-particle system, and that could be explicitly defined in terms of its mechanical state. This, however, is less clear for Maxwell than for Boltzmann.

Gradually, probability was emancipated from this mechanical background. Some isolated papers of Boltzmann [1871b] and Maxwell [1879] already pursued the idea that probabilities characterize an *ensemble* of many many-particle systems rather than a single system. Gibbs's 1902 book adopted this as a uniform coherent viewpoint. However, this ensemble interpretation is still sufficiently vague to be susceptible to different readings. A subjective view of ensembles, closely related to the classical interpretation of Bernoulli and Laplace, has emerged in the 1950s in the work of Jaynes. This paper, will omit a further discussion of this approach. I refer to [Jaynes, 1983; Uffink, 1995; 1996; Balian, 2005] for more details.

3.2 *From Bernoulli to Maxwell (1860)*

The kinetic theory of gases (sometimes called: dynamical theory of gases) is commonly traced back to a passage in Daniel Bernoulli's *Hydrodynamica* of 1738. Previous authors were, of course, quite familiar with the view that gases are composed of a large but finite number of microscopic particles. Yet they usually explained the phenomenon of gas pressure by a static model, assuming repulsive forces between these particles.

Bernoulli's discussion is the first to explain pressure as being due to their motion. He considered a gas as consisting of a great number of particles, moving hither and thither through empty space, and exerting pressure by their incessant collisions

on the walls. With this model, Bernoulli was able to obtain the ideal gas law $pV = \text{const.}$ at constant temperature, predicted corrections to this law at high densities, and argued that the temperature could be taken as proportional to the square of the velocity of the particles. Despite this initial success, no further results were obtained in kinetic gas theory during the next century. By contrast, the view that modeled a gas as a continuum proved much more fertile, since it allowed the use of powerful tools of calculus. Indeed, the few works in the kinetic theory in the early nineteenth century e.g. by Waterston and Herapath were almost entirely ignored by their contemporaries (cf. [Brush, 1976]).

Nevertheless, the kinetic view was revived in the 1850s, in works by Kronig and Clausius. The main stimulus for this revival was the Joule-Mayer principle of the equivalence of heat and work, which led to the First Law of thermodynamics, and made it seem more plausible that heat itself was just a form of motion of gas particles. (A point well-captured in the title of Clausius' 1857 paper: "The kind of motion we call heat", subsequently adopted by Stephen Brush 1976 for his work on the history of this period.)

Clausius also recognized the importance of mutual collisions between the particles of the gas, in order to explain the relative slowness of diffusion when compared with the enormous speed of the particles (estimated at values of 400 m/s or more at ordinary room temperature). Indeed, he argued that in spite of their great speed, the mean free path, i.e. the distance a particle typically travels between two collision, could be quite small (of the order of micrometers) so that the mean displacement per second of particles is accordingly much smaller.

However, Clausius did not pay much attention to the consideration that such collisions would also change the magnitude of the velocities. Indeed, although his work sometimes mentions phrases like "mean speed" or "laws of probability" he does not specify a precise averaging procedure or probability assumption, and his calculations often proceed by crude simplifications (e.g. assuming that all but one of the particles are at rest).

Maxwell (1860)

It was Maxwell's paper of 1860 that really marks the re-birth of kinetic theory. Maxwell realized that if a gas consists of a great number N of moving particles, their velocities will suffer incessant change due to mutual collisions, and that a gas in a stationary state should therefore consist of a mixture of slower and faster particles. More importantly, for Maxwell this was not just an annoying complication to be replaced by simplifying assumptions, but the very feature that deserved further study.

He thus posed the question

Prop. IV. To find the average number of particles whose velocities lie between given limits, after a great number of collisions among a great number of equal particles. [Maxwell, 1860, p. 380].

Denoting this desired average number as $Nf(\vec{v})d^3\vec{v}$, he found a solution to this problem by imposing two assumptions: the distribution function $f(\vec{v})$ should (i) factorize into functions of the orthogonal components of velocity, i.e. there exists some function g such that:

$$(24) \quad f(\vec{v}) = g(v_x)g(v_y)g(v_z),$$

and (ii) be spherically symmetric, i.e.,

$$(25) \quad f(\vec{v}) \text{ depends only on } v = \|\vec{v}\|.$$

He observed that these functional equations can only be satisfied if

$$(26) \quad f(\vec{v}) = Ae^{-v^2/B},$$

where the constant A is determined by normalization: $A = (B\pi)^{-3/2}$; and constant B is determined by relating the mean squared velocity to the absolute temperature — i.e., adopting modern notation: $\frac{3}{2}kT = \frac{m}{2}\langle v^2 \rangle$ — to obtain:

$$(27) \quad f(\vec{v}) = \left(\frac{m}{2\pi kT}\right)^{3/2} e^{-mv^2/2kT}.$$

Maxwell's result led to some novel and unexpected predictions, the most striking being that the viscosity of a gas should be independent of its density, which was, nevertheless, subsequently experimentally verified. Another famous prediction of Maxwell was that in this model the ratio of the specific heats $\gamma = \frac{c_V}{c_p}$ must take the value of $\frac{4}{3}$. This did not agree with the experimentally obtained value of $\gamma = 1.408$.¹⁰

Maxwell's paper is the first to characterize the state of a gas by a distribution function f . It is also the first to call $f(\vec{v})d^3\vec{v}$ a *probability*. Clearly, Maxwell adopted a frequency interpretation of probability. The probability for the velocity to lie within a certain range $d^3\vec{v}$ is nothing but the relative number of particles in the gas with a velocity in that range. It refers to an objective, mechanical property of the gas system, and not to our opinions.

Now an obvious problem with this view is that if the gas contains a finite number of particles, the distribution of velocities must necessarily be discrete, i.e., in Dirac delta notation:

$$(28) \quad f(\vec{v}) = \frac{1}{N} \sum_{i=1}^N \delta(\vec{v} - \vec{v}_i),$$

¹⁰More generally, $c_V/c_p = (f+2)/f$ where f is the number of degrees of freedom of a molecule. This so-called c_V/c_p anomaly haunted gas theory for another half century. The experimental value around 1.4 is partly due to the circumstance that most ordinary gases have diatomic molecules for which, classically, $f = 6$. Quantum theory is needed to explain that one of these degrees is “frozen” at room temperature. Experimental agreement with Maxwell's prediction was first obtained by Kundt and Warburg in 1875 for mercury vapour. (For more details, see [Brush, 1976, p. 353–356]).

and if the energy of the gas is finite and fixed, the distribution should have a bounded support. The function (26) has neither of these properties.

It is not clear how Maxwell would have responded to such problems. It seems plausible that he would have seen the function (26) as representing only a good enough approximation,¹¹ in some sense, to the actual state of the gas but not to be taken too literally, just like actual frequencies in a chance experiment never match completely with their expected values. This is suggested by Maxwell's own illustration of the continuous distribution function as a discrete cloud of points, each of which representing the endpoint of a velocity vector (cf. Fig. 1 from [Maxwell, 1875]). This suggests he thought of an actual distribution more along the lines of (28) than (26). But this leaves the question open in what sense the Maxwell distribution approximates the actual distribution of velocities.

One option here would be to put more emphasis on the phrase "average" in the above quote from Maxwell. That is, maybe f is not intended to represent an actual distribution of velocities but an averaged one. But then, what kind of average? Since an average over the particles has already been performed, the only reasonable options could be an average over time or averaging over an ensemble of similar gas systems. But I can find no evidence that Maxwell conceived of such procedures in this paper. Perhaps one might argue that the distribution (26) is intended as an expectation, i.e. that it represents a reasonable mind's guess about the number of particles with a certain velocity. But in that case, Maxwell's interpretation of probability ultimately becomes classical.

However this may be, it is remarkable that the kinetic theory was thus able to make progress beyond Bernoulli's work by importing mathematical methods (functional equations) involving the representation of a state by continuous functions; though at the price of making this state concept more abstractly connected to physical reality.

A more pressing problem is that the assumptions (24, 25) Maxwell used to derive the form of his distribution do not sit well with its intended frequency interpretation. They seem to reflect a priori desiderata of symmetry, and are perhaps motivated by an appeal to some form of the principle of insufficient reason, in the sense that if there is, in our knowledge, no reason to expect a dependence between the various orthogonal components of velocity, we are entitled to assume they are independent.

This reading of Maxwell's motivations is suggested by the fact that in 1867 he described his 1860 assumption (24) as "the assumption that the probability of a molecule having a velocity resolved parallel to x lying between given limits is not in any way affected by the *knowledge* that the molecule has a given velocity resolved parallel to y " [Maxwell, 1867, emphasis added].

It has been pointed out (see e.g. [Brush, 1976, Vol. II, pp. 183–188]) that Maxwell's 1860 argument seems to have been heavily inspired by [Herschel, 1850] review

¹¹This view was also expressed by [Boltzmann, 1896b]. He wrote, for example: "For a finite number of molecules the Maxwell distribution can never be realized exactly, but only as a good approximation" [Boltzmann, 1909, III, p. 569].




Diagram of Velocities.

Figure 1. An illustration of the Maxwell distribution from [Maxwell, 1875]. Every dot represents the end-point of a velocity vector.

of Quetelet's work on probability. This review essay contained a strikingly similar argument, applied to a marksman shooting at a target, in order to determine the probability that a bullet will land at some distance from the target. What is more, Herschel's essay is firmly committed to the classical interpretation of probability and gives the principle of insufficient reason a central role. Indeed, he explains the (analogue of) condition (25) as "nothing more than the expression of our state of *complete* ignorance of the causes of the errors [i.e. the deviation from the target] and their mode of action" [Herschel, 1850, p. 398, emphasis in the original]. If Maxwell indeed borrowed so much from Herschel, it seems likely that he would also have approved of, or at least be inspired by, this motivation of condition (25).¹²

Whatever may have been Maxwell's original motivation for these assumptions, their dubious nature is also clear from the fact that, in spite of his formulation of the problem (i.e. to determine the form of the function f "after a great number of collisions"), they do not refer to collisions at all. Indeed, it would seem that any motivation for their validity would just as well apply to a gas model consisting of non-colliding (e.g. perfectly transparent) particles as well. As Maxwell himself later remarked about certain derivations in the works of others, one might say that the condition "after a great number of collisions" is intended "rather for the sake of enabling the reader to form a mental image of the material system than as a condition for the demonstration" ([Garber *et al.*, 1995, p. 359]. [Maxwell, 1879]).

3.3 Maxwell (1867)

Whatever the merits and problems of his first paper, Maxwell's next paper on gas theory of 1867 rejected his previous attempt to derive the distribution function from the assumptions (24, 25) as "precarious" and proposed a completely different argument. This time, he considered a model of point particles with equal masses interacting by means of a repulsive central force, proportional to the fifth power of their mutual distance. What is more important, this time the collisions are used in the argument.

Maxwell considers an elastic collision between a pair of particles such that the initial velocities are \vec{v}_1, \vec{v}_2 and final velocities \vec{v}_1', \vec{v}_2' .¹³ These quantities are related by the conservation laws of momentum and energy, yielding four equations, and two parameters depending on the geometrical factors of the collision process.

It is convenient to consider a coordinate frame such that particle 1 is at rest in the origin, and the relative velocity $\vec{v}_2 - \vec{v}_1$ is directed along the negative z axis, and

¹²It is interesting to note that Herschel's review prompted an early and biting criticism of the principle of insufficient reason as applied to frequencies of events by Leslie Ellis, containing the famous observation: "Mere ignorance is no ground for any inference whatsoever. *Ex nihilo nihil*. It cannot be that because we are ignorant of the matter, we know something about it" [Ellis, 1850]. It is not certain, however, whether Maxwell knew of this critique.

¹³In view of the infinite range of the interaction, 'initial' and 'final' are to be understood in an asymptotic sense, i.e. in the limits $t \rightarrow \pm\infty$. An alternative followed in the text is to replace Maxwell's (1867) model with the hard spheres he had considered in 1860.

to use cylindrical coordinates. If (b, ϕ, z) denote the coordinates of the trajectory of the centre of particle 2, we then have $b = const.$, $\phi = const.$, $z(t) = z_0 - \|\vec{v}_2 - \vec{v}_1\|t$ before the collision. In the case where the particles are elastic hard spheres, a collision will take place only if the impact parameter b is less than the diameter d of the spheres. The velocities after the collision are then determined by $\|\vec{v}_1 - \vec{v}_2\|, b$ and ϕ . Transformed back to the laboratory frame, the final velocities \vec{v}_1', \vec{v}_2' can then be expressed as functions of \vec{v}_1, \vec{v}_2, b and ϕ .

Maxwell now assumes what the Ehrenfests later called the *Stoßzahlansatz*: the number of collisions during a time dt , say $N(\vec{v}_1, \vec{v}_2)$, in which the initial velocities \vec{v}_1, \vec{v}_2 within an element $d^3\vec{v}_1 d^3\vec{v}_2$ are changed into final velocities \vec{v}_1', \vec{v}_2' in an element $d^3\vec{v}_1' d^3\vec{v}_2'$ within a spatial volume element $dV = bdbd\phi dz = \|\vec{v}_1 - \vec{v}_2\| bdbd\phi dt$ is proportional to the product of the number of particles with velocity \vec{v}_1 within $d^3\vec{v}_1$ (i.e. $Nf(\vec{v}_1)d^3\vec{v}_1$), and those with velocity \vec{v}_2 within $d^3\vec{v}_2$ (i.e. $Nf(\vec{v}_2)d^3\vec{v}_2$), and that spatial volume element. Thus:

$$(29) \quad N(\vec{v}_1, \vec{v}_2) = N^2 f(\vec{v}_1) f(\vec{v}_2) \|\vec{v}_2 - \vec{v}_1\| d^3\vec{v}_1 d^3\vec{v}_2 bdbd\phi dt.$$

Due to the time reversal invariance of the collision laws, a similar consideration applies to the so-called inverse collisions, in which initial velocities \vec{v}_1', \vec{v}_2' and final velocities \vec{v}_1 and \vec{v}_2 are interchanged. Their number is proportional to

$$(30) \quad N(\vec{v}_1', \vec{v}_2') = N^2 f(\vec{v}_1') f(\vec{v}_2') \|\vec{v}_2' - \vec{v}_1'\| d^3\vec{v}_1' d^3\vec{v}_2' bdbd\phi dt$$

Maxwell argues that the distribution of velocities will remain stationary, i.e. unaltered in the course of time, if the number of collisions of these two kinds are equal, i.e. if

$$(31) \quad N(\vec{v}_1', \vec{v}_2') = N(\vec{v}_1, \vec{v}_2).$$

Moreover, the collision laws entail that $\|\vec{v}_2' - \vec{v}_1'\| = \|\vec{v}_2 - \vec{v}_1\|$ and $d^3\vec{v}_1' d^3\vec{v}_2' = d^3\vec{v}_1 d^3\vec{v}_2$. Hence, the condition (31) may be simplified to

$$(32) \quad f(\vec{v}_1) f(\vec{v}_2) = f(\vec{v}_1') f(\vec{v}_2'), \quad \text{for all } \vec{v}_1, \vec{v}_2.$$

This is the case for the Maxwellian distribution (26). Therefore, Maxwell says, the distribution (26) is a “possible” form.

He goes on to claim that it is also the *only* possible form for a stationary distribution. This claim, i.e. that stationarity of the distribution f can only arise under (32) is nowadays also called the principle of *detailed balancing* (cf. [Tolman, 1938, p. 165]).¹⁴ Although his argument is rather brief, the idea seems to be that for a distribution violating (32), there must (because of the *Stoßzahlansatz*) be two velocity pairs¹⁵ \vec{v}_1, \vec{v}_2 and \vec{u}_1, \vec{u}_2 , satisfying $\vec{v}_1 + \vec{v}_2 = \vec{u}_1 + \vec{u}_2$ and $v_1^2 + v_2^2 = u_1^2 + u_2^2$, such that the collisions would predominantly transform $(\vec{v}_1, \vec{v}_2) \longrightarrow (\vec{u}_1, \vec{u}_2)$ rather

¹⁴The reader might be warned, however, that the name ‘detailed balancing’ is also used to cover somewhat different ideas than expressed here [Tolman, 1938, p. 521].

¹⁵Actually, Maxwell, discusses only velocities of a single molecule. For clarity, I have transposed his argument to a discussion of pairs.

than $(\vec{u}_1, \vec{u}_2) \longrightarrow (\vec{v}_1, \vec{v}_2)$. But then, since the distribution is stationary, there must be a third pair of velocities, (\vec{w}_1, \vec{w}_2) , satisfying similar relations, for which the collisions predominantly produce transitions $(\vec{u}_1, \vec{u}_2) \longrightarrow (\vec{w}_1, \vec{w}_2)$, etc. Now, the distribution can only remain stationary if any such sequence closes into a cycle. Hence there would be cycles of velocity pairs $(\vec{v}_1, \vec{v}_2) \longrightarrow (\vec{u}_1, \vec{u}_2) \longrightarrow \dots \longrightarrow (\vec{v}_1, \vec{v}_2)$ which the colliding particles go through, eventually returning to their original velocities.

Maxwell then argues: “Now it is impossible to assign a reason why the successive velocities of a molecule should be arranged in this cycle rather than in the reverse order” [Maxwell, 1867, p.45]. Therefore, he argues, these two cycles should be equally probable, and, hence, a collision cycle of the type $(\vec{v}_1, \vec{v}_2) \longrightarrow (\vec{v}'_1, \vec{v}'_2)$ is already equally probable as a collision cycle of the type $(\vec{v}'_1, \vec{v}'_2) \longrightarrow (\vec{v}_1, \vec{v}_2)$, i.e. condition (32) holds.

Comments. First, a clear advantage of Maxwell’s 1867 derivation of the distribution function (26) is that the collisions play a crucial role. The argument would not apply if there were no collisions between molecules. A second point to note is that the distribution (26) is singled out because of its *stationarity*, instead of its spherical symmetry and factorization properties. This is also a major improvement upon his previous paper, since stationarity is essential to thermal equilibrium.

A crucial element in the argument is still an assumption about independence. But now, in the *Stoßzahlansatz*, the initial velocities of colliding particles are assumed independent, instead of the orthogonal velocity components of a single particle. Maxwell does not expand on why we should assume this *ansatz*; he clearly regarded it as obvious. Yet it seems plausible to argue that he must have had in the back of his mind some version of the principle of insufficient reason, i.e., that we are entitled to treat the initial velocities of two colliding particles as independent because we have no reason to assume otherwise. Although still an argument from insufficient reason, this is at least a much more plausible application than in the 1860 paper.

A main defect of the paper is his sketchy claim that the Maxwell distribution (26) would be the unique stationary distribution. This claim may be broken in two parts: (a) the cycle argument just discussed, leading Maxwell to argue for detailed balancing; and (b) the claim that the Maxwell distribution is uniquely compatible with this condition.

A demonstration for part (b) was not provided by Maxwell at all; but this gap was soon bridged by Boltzmann (1868) — and Maxwell gave Boltzmann due credit for this proof. But part (a) is more interesting. We have seen that Maxwell here explicitly relied on reasoning from insufficient reason. He was criticized on this point by [Boltzmann, 1872] and also by [Guthrie, 1874].

Boltzmann argued that Maxwell was guilty of begging the question. If we suppose that the two cycles did not occur equally often, then this supposition by itself would provide a reason for assigning unequal probabilities to the two types

of collisions.¹⁶ This argument by Boltzmann indicates, at least in my opinion that he was much less prepared than Maxwell to argue in terms of insufficient reason. Indeed, as we shall see in Section 4, his view on probability seems much more thoroughly frequentist than Maxwell.

In fact Boltzmann later repeatedly mentioned the counterexample of a gas in which all particles are lined up so that they only collide centrally, and move perpendicularly between parallel walls [Boltzmann, 1872 (Boltzmann, 1909, I p. 358); Boltzmann, 1878 (Boltzmann, 1909, II p. 285)]. In this case, the velocity distribution

$$(33) \quad \frac{1}{2} (\delta(v - v_0) + \delta(v + v_0))$$

is stationary too.

Some final remarks on Maxwell's work: As we have seen, it is not easy to pinpoint Maxwell's interpretation of probability. In his (1860), he identifies the probability of a particular molecular state with the relative number of particles that possess this state.¹⁷ Yet, we have also seen that he relates probability to a state of knowledge. Thus, his position may be characterized as somewhere between the classical and the frequentist view.

Note that Maxwell never made any attempt to reproduce the second law. Rather he seems to have been content with the statistical description of thermal equilibrium in gases.¹⁸ All his writings after 1867 indicate that he was convinced that a derivation of the Second Law from mechanical principles was impossible. Indeed, his remarks on the Second Law generally point to the view that the Second Law "has only statistical certainty" (letter to Tait, undated; [Garber *et al.*, 1995, p. 180]), and that statistical considerations were foreign to the principles of mechanics. Indeed, Maxwell was quite amused to see Boltzmann and Clausius engage in a dispute about who had been the first to reduce the Second Law of thermodynamics to mechanics:

It is rare sport to see those learned Germans contending the priority of the discovery that the 2nd law of $\theta\Delta cs$ is the 'Hamiltonsche Prinzip',

¹⁶More precisely, Boltzmann argued as follows: "In order to prove the impossibility [of the hypothesis] that the velocity of [a pair of] molecule[s] changes more often from $[(\vec{v}_1, \vec{v}_2)$ to $(\vec{v}'_1, \vec{v}'_2)]$ than the converse, Maxwell says that there should then exist a closed series of velocities that would be traversed rather in one order than the other. This, however, could not be, he claims, because one could not indicate a reason, why molecules would rather traverse the cycle in one order than the other. But it appears to me that this last claim already presupposes as proven what is to be proved. Indeed, if we assume as already proven that the velocities change as often from (\vec{v}_1, \vec{v}_2) to (\vec{v}'_1, \vec{v}'_2) as conversely, then of course there is no reason why the cycle should rather be run through in one order than the other. But if we assume that the statement to be proven is not yet proved, then the very fact that the velocities of the molecules prefer to change rather from (\vec{v}_1, \vec{v}_2) to (\vec{u}_1, \vec{u}_2) than conversely, rather from (\vec{u}_1, \vec{u}_2) to (\vec{w}_1, \vec{w}_2) than conversely, etc. would provide the reason why the cycle is traversed rather one way than the other" [Boltzmann, 1909, I, p. 319].

¹⁷Curiously, this terminology is completely absent in his 1867 paper.

¹⁸Apart from a rather lame argument in [Maxwell, 1860] analyzed by [Brush, 1976, p.344].

[...] The Hamiltonsche Prinzip, the while, soars along in a region un-
vexed by statistical considerations, while the German Icarus flap their
waxen wings in *nephelococcygia*¹⁹ amid those cloudy forms which the
ignorance and finitude of human science have invested with the incom-
murable attributes of the invisible Queen of Heaven (letter to Tait,
1873; [Garber *et al.*, 1995, p. 225])

Clearly, Maxwell saw a derivation of the Second Law from pure mechanics, “un-
vexed by statistical considerations”, as an illusion. This point appears even more
vividly in his thought experiment of the “Maxwell demon”, by which he showed
how the laws of mechanics could be exploited to produce a violation of the Sec-
ond Law. For an entry in the extensive literature on Maxwell’s demon, I refer
to [Earman and Norton, 1998; 1999; Leff and Rex, 2003; Bennett, 2003; Norton,
2005].

But neither did Maxwell make any effort to reproduce the Second Law on a
unified statistical/mechanical basis. Indeed, the scanty comments he made on
the topic (e.g. in [Maxwell, 1873; Maxwell, 1878b]) rather seem to point in an-
other direction. He distinguishes between what he calls the ‘statistical method’
and the ‘historical’ or ‘dynamical’ (or sometimes ‘kinetic’) method. These are two
modes of description for the same system. But rather than unifying them, Max-
well suggests they are competing, or even incompatible — one is tempted to say
“complementary”-methods, and that it depends on our own knowledge, abilities,
and interests which of the two is appropriate. For example:

In dealing with masses of matter, while we do not perceive the individ-
ual molecules, we are *compelled* to adopt what I have described as the
statistical method, and to abandon the strict dynamical method, in
which we follow every motion by the calculus [Maxwell, 1872, p. 309,
emphasis added].

In this respect, his position stands in sharp contrast to that of Boltzmann, who
made the project of finding this unified basis his lifework.

4 BOLTZMANN²⁰

4.1 *Early work: Stoßzahlansatz and ergodic hypothesis*

Boltzmann had already been considering the problem of finding a mechanical
derivation of the Second Law in a paper of 1866. At that time, he did not know
of Maxwell’s work. But in 1868, he had read both Maxwell’s papers of 1860 and
1867. Like Maxwell, he focuses on the study of gases in thermal equilibrium, in-
stead of the Second Law. He also adopts Maxwell’s idea of characterizing thermal
equilibrium by a probability distribution, and the *Stoßzahlansatz* as the central

¹⁹‘Cloudcuckooland’, an illusory place in Aristophanes’ *The Birds*.

²⁰Parts of this section were also published in [Uffink, 2004].

dynamical assumption. But along the way in this extensive paper, Boltzmann comes to introduce an entirely different alternative approach, relying on what we now call the ergodic hypothesis.

As we saw in section 3.3, Maxwell had derived his equilibrium distribution for two special gas models (i.e. a hard sphere gas in 1860 and a model of point particles with a central r^5 repulsive force acting between them in 1867). He had noticed that the distribution, once attained, will remain stationary in time (when the gas remains isolated), and also argued (but not very convincingly) that it was the *only* such stationary distribution.

In the first section of his 1868a, Boltzmann aims to reproduce and improve these results for a system of an infinite number of hard discs moving in a plane. He regards it as obvious that the equilibrium distribution should be independent of the position of the discs, and that every direction of their velocities is equally probable. It is therefore sufficient to consider the probability distribution over the various values of the velocity $v = \|\vec{v}\|$. However, Boltzmann started out with a somewhat different interpretation of probability in mind than Maxwell. He introduced the probability distribution as follows:

Let $\phi(v)dv$ be the sum of all the instants of time during which the velocity of a disc in the course of a very long time lies between v and $v + dv$, and let N be the number of discs which on average are located in a unit surface area, then

$$(34) \quad N\phi(v)dv$$

is the number of discs per unit surface whose velocities lie between v and $v + dv$ [Boltzmann, 1909, I, p. 50].²¹

Thus, $\phi(v)dv$ is introduced as the relative *time* during which a (given) disc has a particular velocity. But, in the same breath, this is identified with the relative *number* of discs with this velocity.

This remarkable quote shows how he identified two different meanings for the same function. We shall see that this equivocation returned in different guises again and again in Boltzmann's writings.²² Indeed, it is, I believe, the very heart of the ergodic problem, put forward so prominently by the Ehrenfests (cf. paragraph 6.1). Either way, of course, whether we average over time or particles, probabilities are defined here in strictly mechanical terms, and are therefore objective properties of the gas.

²¹Here and below, "Abh." refers to the three volumes of Boltzmann's collected scientific papers [Boltzmann, 1909].

²²This is not to say that he always conflated these two interpretations of probability. Some papers employ a clear and consistent choice for one interpretation only. But then that choice differs between papers, or even in different sections of a single paper. In fact, in [Boltzmann, 1871c] he even multiplied probabilities with different interpretations into one equation to obtain a joint probability. But then in 1872 he conflates them again. Even in his last paper [Boltzmann and Nabl, 1904], Boltzmann identifies two meanings of probability with a simple-minded argument.

Next he goes into a detailed mechanical description of a two-disc collision process. If the variables which specify the initial velocities of two discs before the collision lie within a given infinitesimal range, Boltzmann determines how the collision will transform the initial values of these variables (\vec{v}_i, \vec{v}_j) into the final values (\vec{v}'_i, \vec{v}'_j) in another range. At this point a two-dimensional analogy of the *Stoßzahlansatz* is introduced to obtain the number of collisions per unit of time. As in Maxwell's treatment, this amounts to assuming that the number of such collisions is proportional to the product $\phi(v_1)\phi(v_2)$. In fact:

$$(35) \quad N(\vec{v}_1, \vec{v}_2) \propto N^2 \frac{\phi(v_1)}{v_1} \frac{\phi(v_2)}{v_2} \|\vec{v}_2 - \vec{v}_1\| dv_1 dv_2 dt$$

where the proportionality constant depends on the geometry of the collision.

He observes that if, for all velocities v_i, v_j and all pairs of discs i, j , the collisions that transform the values of the velocities (v_i, v_j) from a first range $dv_i dv_j$ into values v'_i, v'_j within the range $dv'_i dv'_j$ occur equally often as conversely (i.e., equally often as those collisions that transform initial velocities v'_i, v'_j within $dv'_i dv'_j$ into final values v_i, v_j within $dv_i dv_j$), the distribution ϕ will remain stationary. He states "This distribution is therefore the desired one" [Boltzmann, 1909, I p. 55]. Actually, this is the first occasion in the paper at which the desideratum of stationarity of the probability distribution is mentioned.

Using the two-dimensional version of the *Stoßzahlansatz* this desideratum leads to

$$(36) \quad \frac{\phi(v_i)}{v_i} \frac{\phi(v_j)}{v_j} = \frac{\phi(v'_i)}{v'_i} \frac{\phi(v'_j)}{v'_j}$$

He shows [Boltzmann, 1909, p. 57] that the only function obeying condition (36) for all choices of $v_1, v_2, v'_1 v'_2$, compatible with the energy equation $v_1^2 + v_2^2 = v_1'^2 + v_2'^2$, is of the form

$$(37) \quad \phi(v) = 2hve^{-hv^2},$$

for some constant h . Putting $f(v) := v\phi(v)$ we thus obtain the two-dimensional version of the Maxwell distribution (26). Boltzmann does not address the issue of whether the condition (36) is necessary for the stationarity of ϕ .

In the next subsections of 1868a, Boltzmann repeats the derivation, each time in slightly different settings. First, he goes over to the three-dimensional version of the problem, assuming a system of hard spheres, and supposes that one special sphere is accelerated by an external potential $V(\vec{x})$. He shows that if the velocities of all other spheres are distributed according to the Maxwellian distribution (26), the probability distribution of finding the special sphere at place \vec{x} and velocity \vec{v} is $f(\vec{v}, \vec{x}) \propto e^{-h(\frac{1}{2}mv^2 + V(\vec{x}))}$ [Boltzmann, 1909, I, p.63]. In a subsequent subsection, he replaces the spheres by material points with a short-range interaction potential and reaches a similar result.

At this point, (the end of Section I of the [1868a] paper), the argument suddenly switches course. Instead of continuing in this fashion, Boltzmann announces

[Boltzmann, 1909, p. 80] that all the cases treated, and others yet untreated, follow from a much more general theorem. This theorem, which, as we shall see relies on the ergodic hypothesis, is the subject of the second and third Section of the paper. I will limit the discussion to the third section and rely partly on Maxwell's (1879) exposition, which is somewhat simpler and clearer than Boltzmann's own.

The ergodic hypothesis

Consider a general mechanical system of N material points, each with mass m , subject to an arbitrary time-independent potential.²³ In modern notation, let $x = (\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N)$ denote the canonical position coordinates and momenta of the system. Its Hamiltonian is then²⁴

$$(38) \quad H(x) = \frac{1}{2m} \sum_i^N \vec{p}_i^2 + U(\vec{q}_1, \dots, \vec{q}_N).$$

The state x may be represented as a phase point in the mechanical phase space Γ . Under the Hamiltonian equations of motion, this phase point evolves in time, and thus describes a trajectory x_t ($t \in \mathbb{R}$). This trajectory is constrained to lie on a given energy hypersurface $\Gamma_E = \{x \in \Gamma : H(x) = E\}$. Boltzmann asks for the probability (i.e. the fraction of time during a very long period) that the phase point lies in a region $dx = d^3\vec{q}_1 \cdots d^3\vec{p}_N$, which we may write as:

$$(39) \quad \rho(x)dx = \chi(x)\delta(H(x) - E)dx.$$

for some function χ . Boltzmann seems to assume implicitly that this distribution is stationary. This property would of course be guaranteed if the "very long period" were understood as an infinite time. He argues, by Liouville's theorem, that χ is a constant for all points on the energy hypersurface that are "possible", i.e. that are actually traversed by the trajectory. For all other points χ vanishes. If we neglect those latter points, the function χ must be constant over the entire energy hypersurface, and the probability density ρ takes the form

$$(40) \quad \rho_{\text{mc}}(x) = \frac{1}{\omega(E)}\delta(H(x) - E),$$

the micro-canonical distribution, where

$$(41) \quad \omega(E) = \int \delta H(x) = E dx$$

is the so-called structure function.

In particular, one can now form the marginal probability density for the positions $\vec{q}_1, \dots, \vec{q}_N$ by integrating over the momenta:

²³Although Boltzmann does not mention it at this stage, his previous section added the stipulation that the particles are enclosed in a finite space, surrounded by perfectly elastic walls.

²⁴Actually Boltzmann allows the N masses to be different but restricts the potential as being due to external and mutual two-particle forces only, i.e. $H(x) = \sum_i \frac{\vec{p}_i^2}{2m_i} + \sum_{i \leq j} U_{ij}(\|\vec{q}_i - \vec{q}_j\|) + \sum_i U_i(\vec{q}_i)$.

$$(42) \quad \rho_{\text{mc}}(\vec{q}_1, \dots, \vec{q}_N) := \int \rho_{\text{mc}}(x) d^3\vec{p}_1 \cdots d^3\vec{p}_N = \frac{2m}{\omega(E)} \int \delta(\sum \vec{p}_i^2 - 2m(E - U(q))) d\vec{p}_1 \cdots d\vec{p}_N.$$

The integral over the momenta can be evaluated explicitly (it is $2R^{-1}$ times the surface area of a hypersphere with radius $R = \sqrt{2m(E - U)}$ in $n = 3N$ dimensions), to obtain

$$(43) \quad \rho_{\text{mc}}(\vec{q}_1, \dots, \vec{q}_N) = \frac{2m(\pi)^{n/2}}{\omega(E)\Gamma(\frac{n}{2})} (2m(E - U(q)))^{(n-2)/2},$$

where Γ denotes Euler’s gamma function: $\Gamma(x) := \int_0^\infty t^{x-1} e^{-t} dt$.

Similarly, the marginal probability density for finding the first particle with a given momentum component p_{1x} as well as finding the positions of all particles at $\vec{q}_1, \dots, \vec{q}_N$ is

$$(44) \quad \begin{aligned} \rho_{\text{mc}}(p_{1x}, \vec{q}_1, \dots, \vec{q}_N) &= \int \rho_{\text{mc}}(x) dp_{1y} dp_{1z} d^3\vec{p}_2 \cdots d^3\vec{p}_N \\ &= \frac{2m\pi^{(n-1)/2}}{\omega(E)\Gamma(\frac{n-1}{2})} (2m(E - U(q)) - p_{1x}^2)^{(n-3)/2}. \end{aligned}$$

These two results can be conveniently presented in the form of the conditional probability that the x -component of momentum of the first particle has a value between p and $p + dp$, given that the positions have the values $\vec{q}_1 \dots \vec{q}_N$, by taking the ratio of (44) and (43):

$$(45) \quad \rho_{\text{mc}}(p | \vec{q}_1, \dots, \vec{q}_N) dp = \frac{1}{\sqrt{2m\pi}} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})} \frac{(E - U - \frac{p^2}{2m})^{(n-2)/2}}{(E - U)^{(n-3)/2}} dp.$$

This, in essence, is the general theorem Boltzmann had announced. Further, he shows that in the limit where $n \rightarrow \infty$, and the kinetic energy per degree of freedom $\kappa := (E - U)/n$ remains constant, the expression (45) approaches

$$(46) \quad \frac{1}{\sqrt{4\pi m\kappa}} \exp\left(-\frac{p^2}{4m\kappa}\right) dp.$$

This probability thus takes the same form as the Maxwell distribution (26), if one equates $\kappa = \frac{1}{2}kT$. Presumably, it is this result that Boltzmann had in mind when he claimed that all the special cases he has discussed in section 1 of his paper, would follow from the general theorem. One ought to note however, that since U , and therefore κ , depends on the coordinates, the condition $\kappa = \text{constant}$ is different for different values of $(\vec{q}_1, \dots, \vec{q}_n)$.

Some comments on this result.

1. The difference between this approach and that relying on the *Stoßzahlansatz* is rather striking. Instead of concentrating on a gas model in which particles

are assumed to move freely except for their occasional collisions, Boltzmann here assumes a much more general Hamiltonian model with an arbitrary interaction potential $U(\vec{q}_1, \dots, \vec{q}_N)$. Moreover, the probability density ρ is defined over phase space, instead of the space of molecular velocities. This is the first occasion where probability considerations are applied to the state of the mechanical system as whole, instead of its individual particles. If the transition between kinetic gas theory and statistical mechanics may be identified with this caesura, (as argued by the Ehrenfests 1912 and by Klein 1973) it would seem that the transition has already been made right here.

But of course, for Boltzmann the transition did not involve a major conceptual move, thanks to his conception of probability as a relative time. Thus, the probability of a particular state of the total system is still identified with the fraction of time in which that state is occupied by the system. In other words, he had no need for ensembles or non-mechanical probabilistic assumptions.

However, one should note that the equivocation between relative time and relative number of particles, which was comparatively harmless in the first section of the 1868 paper, is now no longer possible in the interpretation of ρ . Consequently, the conditional probability $\rho(p|\vec{q}_1, \dots, \vec{q}_N)$ gives us the relative time that the total system is in a state for which particle 1 has a momentum with x -component between p and $p + dp$, for given values of all the positions. There is no immediate route to conclude that this has anything to do with the relative number of particles with the momentum p . In fact, there is no guarantee that the probability (45) for particle 1 will be the same for other particles too, unless we use the assumption that U is invariant under permutation of the particles. Thus, in spite of their identical form, the distribution (46) has a very different meaning than (26).

2. The transition from (45) to (46), by letting the number of particles become infinite, also seems to be the first instance of a thermodynamical limit. Since the Maxwell distribution is thus recovered only in this limit, Boltzmann's procedure resolves some questions raised above concerning Maxwell's distribution. For a finite number of particles, the distribution (45) always has a finite support, i.e. $\rho_{mc} = 0$ for those values of $p_i^2 \geq 2m(E - U)$. Thus, we do not run into trouble with the finite amount of energy in the gas.

3. Most importantly, the results (45,46) open up a perspective of great generality. It suggests that the probability of the molecular velocities for an isolated system in a stationary state will always assume the Maxwellian form if the number of particles tends to infinity. Notably, this proof seems to completely dispense with any particular assumption about collisions, or other details of the mechanical model involved, apart from the assumption that it is Hamiltonian. Indeed it need not even represent a gas.

4. The main weakness of the present result is its assumption that the trajectory actually visits all points on the energy hypersurface. This is nowadays called the *ergodic hypothesis*.²⁵

²⁵The literature contains some surprising confusion about how the hypothesis got its name. The Ehrenfests borrowed the name from Boltzmann's concept of an "*Ergode*", which he introduced

Boltzmann returned to this issue on the final page of the paper [Boltzmann, 1909, p. 96]. He notes there that there might be exceptions to his theorem, for example, when the trajectory is periodic. However, Boltzmann observed, such cases would be sensitive to the slightest disturbance from outside. They would be destroyed, e.g. by the interaction of a single free atom that happened to be passing by. He argued that these exceptions would thus only provide cases of unstable equilibrium.

Still, Boltzmann must have felt unsatisfied with his own argument. According to an editorial footnote in his collected works [Boltzmann, 1909, I p.96], Boltzmann's personal copy of the paper contains a hand-written remark in the margin stating that the point was still dubious and that it had not been proven that, even in the presence of interaction with a single external atom, the system would traverse all possible values compatible with the energy equation.

Doubts about the ergodic hypothesis

Boltzmann's next paper 1868b was devoted to checking the validity of the ergodic hypothesis in a relatively simple solvable mechanical model. This paper also gives a nice metaphoric formulation of the ergodic hypothesis: if the phase point were a light source, and its motion exceedingly swift, the entire energy surface would appear to us as homogeneously illuminated [Boltzmann, 1909, I, p. 103]. However, his doubts were still not laid to rest. His next paper on gas theory 1871a returns to the study of a detailed mechanical gas model, this time consisting of polyatomic

in [Boltzmann, 1884] and also discussed in his Lectures on Gas Theory [Boltzmann, 1898]. But what did Boltzmann actually understand by an *Ergode*? Brush points out in his translation of [Boltzmann, 1898, p. 297], and similarly in [Brush, 1976, p. 364], that Boltzmann used the name to denote a stationary ensemble, characterized by the microcanonical distribution in phase space. In other words, in the context of Boltzmann's 1898 an *Ergode* is just an microcanonical ensemble, and seems to have nothing to do with the so-called ergodic hypothesis. Brush criticized the Ehrenfests for causing confusion by their terminology.

However, in his original 1884 introduction of the phrase, the name *Ergode* is used for a stationary ensemble with only a *single* integral of motion, i.e. its total energy. As a consequence, the ensemble is indeed micro-canonical, but, what is more, every member of the ensemble satisfies the hypothesis of traversing every phase point with the given total energy. Indeed, in this context, being an element of an *Ergode* implies satisfaction of this hypothesis. Thus, the Ehrenfests were actually justified in baptizing the hypothesis "ergodic".

Another dispute has emerged concerning the etymology of the term. The common opinion, going back at least to the Ehrenfests has been that the word derived from *ergos* (work) and *hodos* (path). [Gallavotti, 1994] has argued however that "undoubtedly" it derives from *ergos* and *eidos* (similar). Now one must grant Gallavotti that it ought to be expected that the etymology of the suffix "-ode" of *ergode* is identical to that of other words Boltzmann coined in this paper, like *Holode*, *Monode*, *Orthode* and *Planode*; and that a reference to path would be somewhat unnatural in these last four cases. However, I don't believe a reference to "eidos" would be more natural. Moreover, it seems to me that if Boltzmann intended this etymology, he would have written "Ergoide" in analogy to planetoid, ellipsoid etc. That he was familiar with this common usage is substantiated by him coining the term "*Momentoide*" for momentum-like degrees of freedom (i.e. those that contribute a quadratic term to the Hamiltonian) in [Boltzmann, 1892]. The argument mentioned by [Cercignani, 1998, p. 141] (that Gallavotti's father is a classicist) fails to convince me in this matter.

molecules, and avoids any reliance on the ergodic hypothesis. And when he did return to the ergodic hypothesis in 1871b, it was with much more caution. Indeed, it is here that he actually first described the worrying assumption as an *hypothesis*, formulated as follows:

The great irregularity of the thermal motion and the multitude of forces that act on a body make it probable that its atoms, due to the motion we call heat, traverse all positions and velocities which are compatible with the principle of [conservation of] energy [Boltzmann, 1909, I p. 284].²⁶

Note that Boltzmann formulates this hypothesis for an arbitrary body, i.e. it is not restricted to gases. He also remarks, at the end of the paper, that “the proof that this hypothesis is fulfilled for thermal bodies, or even is fulfillable, has not been provided” [Boltzmann, I p. 287].

There is a major confusion among modern commentators about the role and status of the ergodic hypothesis in Boltzmann’s thinking. Indeed, the question has often been raised how Boltzmann could ever have believed that a trajectory traverses *all* points on the energy hypersurface, since, as the Ehrenfests conjectured in 1911, and was shown almost immediately in 1913 by Plancherel and Rosenthal, this is mathematically impossible when the energy hypersurface has a dimension larger than 1 (cf. paragraph 6.1).

It is a fact that both his (1868a, Abh. I, p.96) and (1871b, Abh. I, p.284) mention external disturbances as an ingredient in the motivation for the ergodic hypothesis. This might be taken as evidence for ‘interventionism’, i.e. the viewpoint that such external influences are crucial in the explanation of thermal phenomena (cf: [Blatt, 1959; Ridderbos and Redhead, 1998]). Yet even though Boltzmann clearly expressed the thought that these disturbances might help to motivate the ergodic hypothesis, he never took the idea very seriously. The marginal note in the (1868a) paper mentioned above indicated that, even if the system is disturbed, there is still no easy proof of the ergodic hypothesis, and all his further investigations concerning this hypothesis assume a system that is either completely isolated from its environment, or at most acted upon by a static external force. Thus, interventionism did not play a significant role in his thinking.²⁷

It has also been suggested, in view of Boltzmann’s later habit of discretizing continuous variables, that he somehow thought of the energy hypersurface as a discrete manifold containing only finitely many discrete cells [Gallavotti, 1994]. On this reading, obviously, the mathematical no-go theorems of Rosenthal and Plancherel no longer apply. Now it is definitely true that Boltzmann developed a preference towards discretizing continuous variables, and would later apply this

²⁶An equivalent formulation of the ergodic hypothesis is that the Hamiltonian is the only independent integral of the Hamiltonian equations of motion. This version is given in the same paper [Boltzmann, 1909, p. 281-2]

²⁷Indeed, on the rare occasions on which he later mentioned external disturbances, it was only to say that they are “not necessary” [Boltzmann, 1895b]. See also [Boltzmann, 1896, §91].

procedure more and more (although usually adding that this was fictitious and purely for purposes of illustration and more easy understanding, cf. paragraph 4.2). However, there is no evidence in the (1868) and (1871b) papers that Boltzmann implicitly assumed a discrete structure of the mechanical phase space or the energy hypersurface.

Instead, the context of his 1871b makes clear enough how he intended the hypothesis, as has already been argued by [Brush, 1976]. Immediately preceding the section in which the hypothesis is introduced, Boltzmann discusses trajectories for a simple example: a two-dimensional harmonic oscillator with potential $U(x, y) = ax^2 + by^2$. For this system, the configuration point (x, y) moves through the surface of a rectangle. (Cf. Fig. 2. See also [Cercignani, 1998, p. 148].) He then notes that if a/b is rational, (actually: if $\sqrt{a/b}$ is rational) this motion is periodic. However, if this value is irrational, the trajectory will, in the course of time, traverse “*almählich die ganze Fläche*” [Boltzmann, 1909, p. 271] of the rectangle. He says that in this case x and y are *independent*, since for each values of x an infinity of values for y in any interval in its range are possible. The very fact that Boltzmann considers intervals for the values of x and y of arbitrary small sizes, and stressed the distinction between rational and irrational values of the ratio a/b , indicates that he did *not* silently presuppose that phase space was essentially discrete, where those distinctions would make no sense.

Now clearly, in modern language, one should say that if $\sqrt{a/b}$ is irrational the trajectory is *dense* in the rectangle, but not that it traverses all points. Boltzmann did not possess this language. In fact, he could not have been aware of Cantor’s insight that the continuum contains more than a countable infinity of points. Thus, the correct statement that, in the case that $\sqrt{a/b}$ is irrational, the trajectory will traverse, for each value of x , an infinity of values of y within any interval however small, could easily have led him to believe (incorrectly) that *all* values of x and y are traversed in the course of time.

It thus seems eminently plausible, in view of the fact that this discussion immediately precedes the formulation of the ergodic hypothesis, that Boltzmann’s understanding of the ergodic hypothesis is really what Ehrenfests dubbed the *quasi-ergodic hypothesis*: the assumption that the trajectory is dense (i.e. passes arbitrarily close to every point) on the energy hypersurface.²⁸ The quasi-ergodic hypothesis is not mathematically impossible in higher-dimensional phase spaces. However, the quasi-ergodic hypothesis does not entail the desired conclusion that the only stationary probability distribution over the energy surface is micro-canonical.

Nevertheless, Boltzmann remained sceptical about the validity of his hypothesis, and attempted to explore different routes to his goal of characterizing thermal equilibrium in mechanics. Indeed, both the preceding 1871a and his next paper 1871c present alternative arguments, with the explicit recommendation that they avoid hypotheses. In fact, he did not return to the ergodic hypothesis at all until

²⁸Or some hypothesis compatible with the quasi-ergodic hypothesis. As it happens, Boltzmann’s example is also compatible with the measure-theoretical hypothesis of ‘metric transitivity’ (cf. paragraph 6.1).




Figure 2. Trajectories in configuration space for a two-dimensional harmonic oscillator with potential $U(x, y) = ax^2 + by^2$. Illustrating the distinction between (i) the case where $\sqrt{a/b}$ is rational (here $4/7$) and (ii) irrational ($1/e$). Only a fragment of the latter trajectory has been drawn.

the 1880s (stimulated by Maxwell's 1879 review of the last section of Boltzmann's 1868 paper). At that time, perhaps feeling fortified by Maxwell's authority, he was to express much more confidence in the ergodic hypothesis. However, after 1885, this confidence disappears again, and although he mentions the hypothesis occasionally in later papers, he never assumes its validity. Most notably, the ergodic hypothesis is not even mentioned in his *Lectures on Gas Theory* 1896, 1898.

To sum up, what role did the ergodic hypothesis play for Boltzmann? It seems that Boltzmann regarded the ergodic hypothesis as a special dynamical assumption that may or may not be true, depending on the nature of the system, and perhaps also on its initial state and the disturbances from its environment. Its role was simply to help derive a result of great generality: for any system for which the hypothesis is true, its equilibrium state is characterized by (45), from which an analogy to the Maxwell distribution may be recovered in the limit $N \rightarrow \infty$, regardless of any details of the inter-particle interactions, or indeed whether the system represented is a gas, fluid, solid or any other thermal body.

As we discussed in paragraph 1.4, the Ehrenfests 1912 have suggested that the ergodic hypothesis played a much more fundamental role. In particular, if the

hypothesis is true, averaging over an (infinitely) long time would be identical to phase averaging with the microcanonical distribution. Thus, they suggested that Boltzmann relied on the ergodic hypothesis in order to equate time averages and phase averages, or in other words, to equate two meanings of probability (relative time and relative volume in phase space.) There is however *no* evidence that Boltzmann ever followed this line of reasoning neither in the 1870s, nor later. He simply never gave any justification for equivocating time and particle averages, or phase averages, at all. Presumably, he thought nothing much depended on this issue and that it was a matter of taste.

4.2 *The Boltzmann equation and H-theorem (1872)*

In 1872 Boltzmann published one of his most important papers. It contained two celebrated results nowadays known as the Boltzmann equation and the *H*-theorem. The latter result was the basis of Boltzmann's renewed claim to have obtained a general theorem corresponding to the Second Law. This paper has been studied and commented upon by numerous authors, and an entire translation of the text has been provided by [Brush, 1966]. Thus, for the present purposes, a succinct summary of the main points might have been sufficient. However, there is still dispute among modern commentators about its actual content.

The issue at stake in this dispute is the question whether the results obtained in this paper are presented as necessary consequences of the mechanical equations of motion, or whether Boltzmann explicitly acknowledged that they would allow for exceptions. Klein has written:

I can find no indication in his 1872 memoir that Boltzmann conceived of possible exceptions to the *H*-theorem, as he later called it [Klein, 1973, p. 73].

Klein argues that Boltzmann only came to acknowledge the existence of such exceptions thanks to Loschmidt's critique in 1877. An opposite opinion is expressed by [von Plato, 1994]. Calling Klein's view a "popular image", he argues that, already in 1872, Boltzmann was well aware that his *H*-theorem had exceptions, and thus "already had a full hand against his future critics". Indeed, von Plato states that

Contrary to a widely held opinion, Boltzmann is not in 1872 claiming that the Second Law and the Maxwellian distribution are *necessary* consequences of kinetic theory [von Plato, 1994, p. 81].

So it might be of some interest to try and settle this dispute.

Boltzmann (1872) starts with an appraisal of the role of probability theory in the context of gas theory. The number of particles in a gas is so enormous, and their movements are so swift that we can observe nothing but average values. The determination of averages is the province of probability calculus. Therefore, "the problems of the mechanical theory of heat are really problems in probability

calculus" [Boltzmann, 1909, I, p. 317]. But, Boltzmann says, it would be a mistake to believe that the theory of heat would therefore contain uncertainties.

He emphasizes that one should not confuse incompletely proven assertions with rigorously derived theorems of probability theory. The latter are necessary consequences of their premisses, just like in any other theory. They will be confirmed by experience as soon as one has observed a sufficiently large number of cases. This last condition, however, should be no significant problem in the theory of heat because of the enormous number of molecules in macroscopic bodies. Yet, in this context, one has to make doubly sure that we proceed with the utmost rigour.

Thus, the message expressed in the opening pages of this paper seems clear enough: the results Boltzmann is about to derive are advertised as doubly checked and utterly rigorous. Still, they are theoretical. Their relationship with experience might be less secure, since any probability statement is only reproduced in observations by sufficiently large numbers of independent data. Thus, Boltzmann would have allowed for exceptions in the relationship between theory and observation, but not in the relation between premisses and conclusion.

He continues by saying what he means by probability, and repeats its equivocation as a fraction of time and the relative number of particles that we have seen earlier in 1868:

If one wants [...] to build up an exact theory [...] it is before all necessary to determine the probabilities of the various states that one and the same molecule assumes in the course of a very long time, and that occur simultaneously for different molecules. That is, one must calculate how the number of those molecules whose states lie between certain limits relates to the total number of molecules [Boltzmann, 1909, I p. 317].

However, this equivocation is not vicious. For most of the paper the intended meaning of probability is always the relative number of molecules with a particular molecular state. Only at the final stages of his paper [Boltzmann, 1909, I, p. 400] does the time-average interpretation of probability (suddenly) recur.

Boltzmann says that both Maxwell and he had attempted the determination of these probabilities for a gas system but without reaching a complete solution. Yet, on a closer inspection, "it seems not so unlikely that these probabilities can be derived on the basis of the equations of motion alone..." [Boltzmann, 1909, I, p. 317]. Indeed, he announces, he has solved this problem for gases whose molecules consist of an arbitrary number of atoms. His aim is to prove that whatever the initial distribution of state in such a system of gas molecules, it must inevitably approach the distribution characterized by the Maxwellian form (*ibid.* p. 320).

The next section specializes to the simplest case of monatomic gases and also provides a more complete specification of the problem he aims to solve. The gas molecules are contained in a fixed vessel with perfectly elastic walls. They interact with each other only when they approach each other at very small distances. These interactions can be mimicked as collisions between elastic bodies. Indeed, these

bodies are modeled as hard spheres [Boltzmann, 1909, I, p. 320]. Boltzmann represents the state of the gas by a time-dependent distribution function $f_t(\vec{v})$, called the “distribution of state”, which gives us, at each time t , the relative number of molecules with velocity between \vec{v} and $\vec{v} + d^3\vec{v}$.²⁹

He also states two more special assumptions:

1. Already in the initial state of the gas, each direction of velocity is equally probable. That is:

$$(47) \quad f_0(\vec{v}) = f_0(v).$$

It is assumed as obvious that this will also hold for any later time.

2. The gas is spatially uniform within the container. That is, the relative number of molecules with their velocities in any given interval, and their positions in a particular spatial region R does not depend on the location of R in the available volume.

The next and crucial assumption used by Boltzmann to calculate the change in the number of particles with a velocity \vec{v}_1 per unit time, is the *Stoßzahlansatz*, (29) and (30).

For modern readers, there are also a few unstated assumptions that go into the construction of this equation. First, the number of molecules must be large enough so that the (discrete) distribution of their velocities can be well approximated by a continuous and differentiable function f . Secondly, f changes under the effect of binary collisions only. This means that the density of the gas should be low (so that three-particle collisions can be ignored) but not too low (which would make collisions too infrequent to change f at all). These two requirements are already hard enough to put in a mathematically precise form. The modern explicitation is that of taking the so-called Boltzmann-Grad limit (cf. paragraph 6.4). The final (unstated) assumption is that all the above assumptions remain valid in the course of time.

He addresses his aim by constructing a differentio-integral evolution equation for f_t , by taking the difference of (29) and (30) and integrating over all variables except \vec{v}_1 and t . The result (in a modern notation) is the *Boltzmann equation*:

$$(48) \quad \frac{\partial f_t(\vec{v}_1)}{\partial t} = N \int_0^d b db \int_0^{2\pi} d\phi \int_{\mathbb{R}^3} d^3\vec{v}_2 \|\vec{v}_2 - \vec{v}_1\| \left(f_t(\vec{v}_1') f_t(\vec{v}_2') - f_t(\vec{v}_1) f_t(\vec{v}_2) \right)$$

which describes the change of f in the course of time, when this function at some initial time is given. (Recall from paragraph 3.3 that the primed velocities are to be thought of as functions of the unprimed velocities and the geometrical parameters of the collision: $\vec{v}'_i = \vec{v}'_i(\vec{v}_1, \vec{v}_2, b, \phi)$, and d denotes the diameter of the hard spheres.)

²⁹ Actually Boltzmann formulated the discussion in terms of a distribution function over kinetic energy rather than velocity. I have transposed this into the latter, nowadays more common formulation.

The H-theorem

Assuming that the Boltzmann equation (48) is valid for all times, one can prove, after a few well-known manipulations, that the following quantity

$$(49) \quad H[f_t] := \int f_t(\vec{v}) \ln f_t(\vec{v}) d^3\vec{v}$$

decreases monotonically in time, i.e.

$$(50) \quad \frac{dH[f_t]}{dt} \leq 0;$$

as well as its stationarity for the Maxwell distribution, i.e.:

$$(51) \quad \frac{dH[f_t]}{dt} = 0 \quad (\forall t) \quad \text{iff} \quad f_t(v) = Ae^{-Bv^2}.$$

Boltzmann concludes Section I of the paper as follows:

It has thus been rigorously proved that whatever may have been the initial distribution of kinetic energy, in the course of time it must necessarily approach the form found by Maxwell. [...] This [proof] actually gains much in significance because of its applicability to the theory of multi-atomic gas molecules. There too, one can prove for a certain quantity [H] that, because of the molecular motion, this quantity can only decrease or in the limiting case remain constant. Thus, one may prove that because of the atomic movement in systems consisting of arbitrarily many material points, there always exists a quantity which, due to these atomic movements, cannot increase, and this quantity agrees, up to a constant factor, exactly with the value that I found in [[Boltzmann, 1871c]] for the well-known integral $\int dQ/T$.

This provides an analytical proof of the Second Law in a way completely different from those attempted so far. Up till now, one has attempted to proof that $\int dQ/T = 0$ for a reversible (*umkehrbaren*) cyclic³⁰ process, which however does not prove that for an irreversible cyclic process, which is the only one that occurs in nature, it is always negative; the reversible process being merely an idealization, which can be approached more or less but never perfectly. Here, however, we immediately reach the result that $\int dQ/T$ is in general negative and zero only in a limit case... [Boltzmann, 1909, I, p. 345]

Thus, as in his 1866 paper, Boltzmann claims to have a rigorous, analytical and general proof of the Second Law. From our study of the paper until now, (i.e. section I) it appears that Klein's interpretation is more plausible than von Plato's. I postpone a further discussion of this dispute to paragraph 4.2, after a brief look at the other sections of the paper.

³⁰The term "cyclic" is missing in Brush's translation, although the original text does speak of "Kreisprozeß". The special notation \oint for cyclic integrals was not introduced until much later.

Further sections of Boltzmann (1872)

Section II is entitled “Replacement of integrals by sums” and devoted to a repetition of the earlier arguments, now assuming that the kinetic energies of the molecules can only take values in a discrete set $\{0, \epsilon, 2\epsilon, \dots, p\epsilon\}$. Boltzmann shows that in the limit $\epsilon \rightarrow 0$, $p\epsilon \rightarrow \infty$ the same results are recovered.

Many readers have been surprised by this exercise, which seems rather superfluous both from a didactic and a logical point of view. (However, some have felt that it foreshadowed the advent of quantum theory.) Boltzmann offers as motivation for the detour that the discrete approach is clearer than the previous one. He argues that integrals only have a symbolic meaning, as a sum of infinitely many infinitesimal elements, and that a discrete calculation yields more understanding. He does not argue, however, that it is closer to physical reality. Be that as it may, the section does eventually take the limit, and recovers the same results as before.

The third section treats the case where the gas is non-uniform, i.e., when condition 2 above is dropped. For this case, Boltzmann introduces a generalized distribution function $f_t(\vec{r}, \vec{v})$, such that $f_t d^3\vec{r} d^3\vec{v}$ represents the relative number of particles with a position in a volume element $d^3\vec{r}$ around \vec{r} and a velocity in an element $d^3\vec{v}$ around \vec{v} .

He obtains a corresponding generalized Boltzmann equation:

$$(52) \quad \frac{\partial f_t(\vec{r}, \vec{v})}{\partial t} + \vec{v} \cdot \nabla_x f_t + \frac{\vec{F}}{m} \cdot \nabla_v f_t = N \int bdbd\phi d^3\vec{v}_2 \|\vec{v}_2 - \vec{v}_1\| \left(f_t(\vec{r}, \vec{v}'_1) f_t(\vec{r}, \vec{v}'_2) - f_t(\vec{r}, \vec{v}_1) f_t(\vec{r}, \vec{v}_2) \right)$$

where \vec{F} denotes an external force field on the gas. The quantity H now takes the form $H[f_t] := \int f_t(\vec{r}, \vec{v}) d^3\vec{r} d^3\vec{v}$; and a generalization of the H -theorem $dH/dt \leq 0$ is obtained.

The last three sections are devoted to polyatomic molecules, and aim to obtain generalized results for this case too. The key ingredient for doing so is, of course, an appropriately generalized *Stoßzahlansatz*. The formulation of this assumption is essentially the same as the one given in his paper on poly-atomic molecules 1871a, which was later shown wrong and corrected by Lorentz. I will not go into this issue (cf. [Lorentz, 1887; Boltzmann, 1887b; Tolman, 1938]).

An interesting passage occurs at the very end of the paper, where he expands on the relationship between H and entropy. He considers a monatomic gas in equilibrium. The stationary distribution of state is given as:

$$(53) \quad f^*(\vec{r}, \vec{v}) = V^{-1} \left(\frac{3m}{4\pi T} \right)^{3/2} \exp\left(\frac{-3mv^2}{4T}\right)$$

where V is the volume of the container. (Note that in comparison with (27), Boltzmann adopts units for temperature that make $k = 2/3$.) He shows that

$$(54) \quad H[f^*] := \int f^* \log f^* dx dv = -N \log V \left(\frac{4\pi T}{3m} \right)^{3/2} - \frac{3}{2}N;$$

which agrees (assuming $S = -kNH[f^*]$) with the thermodynamical expression for the ideal gas (16) up to an additive constant. A similar result holds for the polyatomic gas.

Remarks and problems

1. The role of probability. As we have seen, the H -theorem formed the basis of a renewed claim by Boltzmann to have obtained a theorem corresponding to the full Second Law (i.e. including both parts) at least for gases. A main difference from his 1866 claim, is that he now strongly emphasizes the role of probability calculus in his derivation. It is clear that the conception of probability expounded here is thoroughly frequentist and that he takes ‘the laws of probability’ as empirical statements. Furthermore, probabilities can be fully expressed in mechanical terms: the probability distribution f is nothing but the relative number of particles whose molecular states lie within certain limits. Thus, there is no conflict between his claims that on the one hand, “the problems of the mechanical theory of heat are really problems in probability calculus” and that the probabilities themselves are derived on the basis of the equations of motion alone, on the other hand. Indeed, it seems to me that Boltzmann’s emphasis on the crucial role of probability in this paper is only intended to convey that probability theory provides a particularly useful and appropriate language for discussing mechanical problems in gas theory. There is no indication in this paper yet that probability theory could play a role by furnishing assumptions of a non-mechanical nature, i.e., independent of the equations of motion (cf. [Boltzmann and Nabl, 1904, p. 520]).

2. The role of the Stoßzahlansatz. Note that Boltzmann stresses the generality, rigour and “analyticity” of his proof. He puts no emphasis on the special assumptions that go into the argument. Indeed, the *Stoßzahlansatz*, later identified as the key assumption that is responsible for the time-asymmetry of the H -theorem, is announced as follows

The determination [of the number of collisions] can only be obtained in a truly tedious manner, by consideration of the relative velocities of both particles. But since this consideration has, apart from its tediousness, not the slightest difficulty, nor any special interest, and because the result is so simple that one might almost say it is self-evident I will only state this result.” [Boltzmann, 1909, I, p. 32].)

It thus seems natural that Boltzmann’s contemporaries must have understood him as claiming that the H -theorem followed necessarily from the dynamics of the mechanical gas model.³¹ I can find no evidence in the paper that he intended this claim to be read with a pinch of salt, as [von Plato, 1991, p. 81] has argued.

³¹Indeed this is *exactly* how Boltzmann’s claims were understood. For example, the recommendation written in 1888 for his membership of the Prussian Academy of Sciences mentions as his main feat that Boltzmann had proven that, whatever its initial state, a gas must necessarily approach the Maxwellian distribution [Kirsten and Körber, 1975, p.109].

Is there then no evidence at all for von Plato's reading of the paper? Von Plato refers to a passage from Section II, where Boltzmann repeats the previous analysis by assuming that energy can take on only discrete values, and replacing all integrals by sums. He recovers, of course, the same conclusion, but now adds a side remark, which touches upon the case of non-uniform gases:

Whatever may have been the initial distribution of states, there is one and only one distribution which will be approached in the course of time. [...] This statement has been proved for the case where the distribution of states was already initially uniform. It must also be valid when this is not the case, i.e. when the molecules are initially distributed in such a way that in the course of time they mix among themselves more and more, so that after a very long time the distribution of states becomes uniform. This will always be the case, with the exception of very special cases, e.g. when all molecules were initially situated along a straight line, and were reflected by the walls onto this line [Boltzmann, 1909, I, p. 358].

It is this last remark that, apparently, led to the view that after all Boltzmann did already conceive of exceptions to his claims. However, I should say that this passage does not convince me. True enough, Boltzmann in the above quote indicates that there are exceptions. But he mentions them only in connection with an *extension* of his results to the case when the gas is not initially uniform, i.e. when condition (2) above is dropped. There can be no doubt that under the assumption of the conditions (1) and (2), Boltzmann claimed the rigorous validity of the *H*-theorem. (Curiously, his more systematic treatment of the non-uniform gas (Section III of 1872) does not mention any exception to the claim that "*H* can only decrease" [Boltzmann, 1909, I p. 362].

As a matter of fact, when Loschmidt formulated the objection, it happened to be by means of an example of a non-uniform gas (although nothing essential depended on this). Thus, if Boltzmann had in 1872 a "full hand against his future critics", as von Plato claims, one would expect his reply to Loschmidt's objection to point out that Loschmidt was correct but that he had already anticipated the objection. Instead, he accused Loschmidt of a fallacy (see paragraph 4.3 below).

But apart from the historical issue of whether Boltzmann did or did not envisage exceptions to his *H*-theorem, it seems more important to ask what kind of justification Boltzmann might have adduced for the *Stoßzahlansatz*. An attempt to answer this question must be somewhat speculative, since, as we have seen, Boltzmann presented the assumption as "almost self-evident" and "having no special interest", and hence presumably as not in need of further explanation. Still the following remarks may be made with some confidence.

First, we have seen that Maxwell's earlier usage of the assumption was never far away from an argument from insufficient reason. Thus, in his approach, one could think of the *Stoßzahlansatz* as expressing that we have no reason to expect any influence or correlation between any pair of particles that are about to collide.

The assumption would then appear as a probabilistic assumption, reflecting a ‘reasonable judgment’, independent from mechanics.

In contrast, Boltzmann’s critique of Maxwell’s approach (cf. footnote 16) suggests that he did not buy this argument for insufficient reason. But since the *Stoßzahlansatz* clearly cannot be conceived of as an assumption about dynamics — like the ergodic hypothesis —, this leaves only the option that it must be due to a special assumption about the mechanical state of the gas. Indeed, in the years 1895–6, when Boltzmann acknowledged the need for the *ansatz* in the proof of his *H*-theorem more explicitly — referring to it as “Assumption A” [Boltzmann, 1895] or “the hypothesis of molecular disorder” [Boltzmann, 1896] —, he formulated it as an assumption *about* the state of the gas.

Yet, even in those years, he would also formulate the hypothesis as expressing that “haphazard governs freely” [Boltzmann, 1895, Abh. III, p. 546] or “that the laws of probability are applicable for finding the number of collisions” [Boltzmann, 1895b]. Similarly, he describes states for which the hypothesis fails as contrived “so as to intentionally violate the laws of probability” [Boltzmann, 1896, §3]. However, I think these quotations should not be read as claims that the *Stoßzahlansatz* was a consequence of probability theory itself. Rather, given Boltzmann’s empirical understanding of “the laws of probability”, they suggest that Boltzmann thought that, as a matter of empirical fact, the assumption would ‘almost always’ hold, even if the gas was initially very far from equilibrium.

3. The *H*-theorem and the Second Law. Note that Boltzmann misconstrues, or perhaps understates, the significance of his results. Both the Boltzmann equation and the *H*-theorem refer to a body of gas in a fixed container that evolves in isolation from its environment. There is no question of heat being exchanged by the gas during a process, let alone in an irreversible cyclic process. His comparison in the quotation on page 965 with Clausius’ integral $\int dQ/T$ (i.e. $\oint dQ/T$ in equation (18) above) is therefore really completely out of place.

The true import of Boltzmann’s results is rather that they provide (i) a generalization of the entropy concept to non-equilibrium states,³² and (ii) a claim that this non-equilibrium entropy $-kH$ increases monotonically as the isolated gas evolves for non-equilibrium towards an equilibrium state. The relationship with the Second Law is, therefore, somewhat indirect: On the one hand, Boltzmann proves much more than was required, since the second law does not speak of non-equilibrium entropy, nor of monotonic increase; on the other hand it proves also less, since Boltzmann does not consider the increase of entropy in general adiabatic processes.

³²Boltzmann emphasized that his expression for entropy should be seen as an *extension* of thermodynamic entropy to non-equilibrium states in [1877b, (Boltzmann, 1909, II, p. 218); 1896, §5]. Of course there is no guarantee that this generalization is the *unique* candidate for a non-equilibrium entropy.

4.3 Boltzmann (1877a): the reversibility objection

According to [Klein, 1973], Boltzmann seemed to have been satisfied with his treatments of 1871 and 1872 and turned his attention to other matters for a couple of years. He did come back to gas theory in 1875 to discuss an extension of the Boltzmann equation to gases subjected to external forces. But this paper does not present any fundamental changes of thought. (However, it does contain some further elucidation, for example, it mentions for the first time that the derivation of the Boltzmann equation requires that the gas is so dilute that collisions between three or more particles simultaneously can be ignored).

However, the 1875 paper did contain a result which, two years later, led to a debate with Loschmidt. Boltzmann showed that (52) implied that a gas in equilibrium in an external force field (such as the earth's gravity) should have the same average kinetic energy at all heights and therefore, a uniform temperature; while its pressure and density would of course vary with height. This conclusion conflicted with the intuition that when molecules travel upwards, they must do work against the gravitational field, and pay for this by having a lower kinetic energy at greater heights.

Now Boltzmann (1875) was not the first to reach the contrary result, and Loschmidt was not the first to challenge it. Maxwell and Guthrie entered into a debate on the very same topic in 1873. But actually their main point of contention need not concern us very much. The discussion between Loschmidt and Boltzmann is particularly important for quite another issue, which Loschmidt only introduced as an side remark. Considering a gas container in a homogeneous gravitational field, Loschmidt discussed a situation where initially all atoms except one lie at rest at the bottom of the container. The single moving atom could then, by collisions, stir the others and send them into motion until a "stationary state", characterized by the Maxwell distribution, is obtained. He continues

By the way, one should be careful about the claim that in a system in which the so-called stationary state has been achieved, starting from an arbitrary initial state, this average state can remain intact for all times. I believe, rather, that one can make this prediction only for a short while with full confidence.

Indeed, if in the above case, after a time τ which is long enough to obtain the stationary state, one suddenly assumes that the velocities of all atoms are reversed, we would obtain an initial state that would appear to have the same character as the stationary state. For a fairly long time this would be appropriate, but gradually the stationary state would deteriorate, and after passage of the time τ we would inevitable return to our initial state: only one atom has absorbed all kinetic energy of the system [...], while all other molecules lie still on the bottom of the container.

Obviously, in every arbitrary system the course of events must be become retrograde when the velocities of all its elements are reversed

[Loschmidt, 1876, p. 139].

Boltzmann's response (1877a)

Boltzmann's response to Loschmidt is somewhat confusing. On the one hand, he acknowledges that Loschmidt's objection is "quite ingenious and of great significance for the correct understanding of the Second Law." However, he also brands the objection as a "fallacy" and a "sophism".³³ But then, two pages later again, the argument is "of the greatest importance since it shows how intimately connected are the Second Law and probability theory."

The gist of the response is this. First, Boltzmann captures the essential core of the problem in an admirably clear fashion:

"Every attempt to prove, from the nature of bodies and the laws of interaction for the forces they exert among each other, without any assumption about initial conditions, that

$$(55) \quad \int \frac{dQ}{T} \leq 0$$

must be in vain" [Boltzmann, 1909, II. p.119–121].

The point raised here is usually known as the *reversibility objection*. And since the H -theorem (which only received this name in the 1890s) was presented in 1872 as a general proof that $\int \frac{dQ}{T} \leq 0$ (cf. the long quotation on page 965), it would imply that this theorem was invalid. Boltzmann aims to show, however, that this objection is a fallacy. His argument might be dissected into 5 central points.

1. Conceding that the proof cannot be given. Boltzmann says that a proof that every distribution must with absolute necessity evolve towards a uniform distribution cannot be given, claiming that this fact "is already taught by probability theory". Indeed, he argues, even a very non-uniform distribution of state is, although improbable to the highest degree, not impossible. Thus, he admits that there are initial states for which H increases, just as well as those for which H decreases. This admission, of course, is hard to rhyme with his professed purpose of showing that it is fallacious to conclude that some assumption about the initial state would be needed.

Note that this passage announces a major conceptual shift. Whereas the 1872 paper treated the distribution of state f_t as if it *defines* probability (i.e. of molecular velocities), this time the distribution of states is itself something which can be to a higher or lesser degree "probable". That is: probabilities are *attributed*

³³The very fact that Boltzmann called this conclusion — which by all means and standards is *correct* — a fallacy shows, in my opinion, that he had not anticipated the objection. In fact, how much Boltzmann had yet to learn from Loschmidt's objection is evident when we compare this judgment to a quotation from his *Lectures on Gas Theory* [1898, p. 442]: "this one-sidedness [of the H -theorem] lies uniquely and solely in the initial conditions."

to distributions of state, i.e. the distribution of state itself is treated as a random variable. This shift in viewpoint became more explicit in his (1877b); as we will discuss in section 4.4 below.

2. Rethinking the meaning of “probability”. Boltzmann argues that every distribution of state, whether uniform or non-uniform, is equally improbable. But there are “infinitely many” more uniform distributions of state than non-uniform distributions. Here we witness another conceptual shift. In (1872), the term “distribution of state” referred to the function $f(\vec{v})$ or $f(\vec{r}, \vec{v})$, representing the relative numbers of molecules with various molecular states. In that sense, there would, of course, only be a *single* uniform distribution of state: the Maxwellian distribution function (53). But since Boltzmann now claims there are many, he apparently uses the term “distribution of state” to denote a much more detailed description, that includes the velocity and position of every individual molecule, so that permutations of the molecules yield a different distribution of state. That is, he uses the term in the sense of what we would nowadays call a microstate, and what he himself would call a “*Komplexion*” a few months later in his (1877b) — on which occasion he would reserve the name ‘distribution of state’ for the macrostate.

Note that Boltzmann assumes every *Komplexion* to be equally probable (or improbable) so that the probability of a particular distribution of state is determined by the relative numbers. Indeed he remarks that it might be interesting to calculate the probabilities of state distributions by determining the ratio of their numbers; this suggestion is also worked out in his subsequent paper of 1877b.

This, indeed, marks another conceptual change. Not only are probabilities attributed to distributions of state instead of being defined by them; they are determined by an equiprobability assumption. Boltzmann does not explicitly motivate the assumption. In view of the discussion in paragraph 3.1, one might conjecture that he must have had something like Laplace’s principle of insufficient reason in mind, which makes any two cases which, according to our information are equally possible, also equally probable. But this would indicate an even larger conceptual change; and not just because Boltzmann is broadly a frequentist concerning probability. Also, the principle of insufficient reason, or any similar assumption, makes sense only from the view point that probability is a non-mechanical notion: it reflects our belief or information about a system. I cannot find any evidence that he accepted this idea. Of course it is also possible to conjecture that he silently fell back upon the ergodic hypothesis. But this conjecture also seems unlikely, given his avoidance of the hypothesis since 1871.

3. A claim about evolutions. Boltzmann says: “Only from the fact that there are so many more uniform than non-uniform distributions of state [i.e.: microstates] follows the larger probability that the distribution will become uniform in the course of time” (p. 120). More explicitly, he continues:

[...] one can prove that infinitely many more initial states evolve after a long time towards a more uniform distribution of states than to a less uniform one, and that even in that latter case, these states will become uniform after an even longer time [Boltzmann, 1909, II, p. 120]³⁴

Note that this is a claim about evolutions of microstates. In fact, it is the first case of what the Ehrenfests later called a *statistical H-theorem*, but what is perhaps better called a *statistical reading* of the *H-theorem*, since in spite of Boltzmann's assertion, no proof is offered.

4. The (im)probability of Loschmidt's initial state. Boltzmann maintains that the initial conditions considered by Loschmidt only have a minute probability. This is because it is obtained by a time evolution and velocity reversal of a non-uniform microstate. Since both time evolution and velocity reversal are one-to-one mappings (or more to the point: they preserve the Liouville measure), these operations should not affect the number or probability of states. Hence, the probability of Loschmidt's state is equal to that of the special non-uniform state from which it is constructed. But by point 2 above, there are infinitely many more uniform states than non-uniform states, so the probability of Loschmidt's state is extraordinarily small.

5. From (im)probability to (im)possibility. The final ingredient of Boltzmann's response is the claim that whatever has an extraordinarily small probability is practically impossible.

The conclusion of Boltzmann's argument, based on these five points, is that the state selected by Loschmidt may be considered as practically impossible. Note that this is a completely static argument; i.e., its logic relies merely on the points 1,2,4 and 5, and makes no assumption about evolutions, apart from the general feature that the dynamical evolution conserves states (or measure). Indeed, point 3, i.e. the statistical reading of the *H-theorem*, is not used in the argument.

As a consequence, the argument, although perfectly consistent, shows more than Boltzmann can possibly have wanted. The same reasoning that implies Loschmidt's initial state can be ignored, also excludes other non-uniform states. In particular, the same probability should be assigned to Loschmidt's initial state *without* the reversal of velocities. But that state *can* be produced in the laboratory, and, presumably, should not be considered as practically impossible. Indeed, if we adopt the rule that all non-uniform states are to be ignored on account of their low probability, we end up with a consideration of uniform states only, i.e. the theory would be reduced to a description of equilibrium, and the *H-theorem* reduced to $dH/dt = 0$, and any time-asymmetry is lost.

This, surely, is too cheap a victory over Loschmidt's objection. What one would like to see in Boltzmann's argument is a greater role for assumptions about the time evolution in order to substantiate his statistical reading of the *H-theorem*.

³⁴The clause about 'the latter case' is absent in the translation by [Brush, 2003, p. 366].

Summing up: From this point on, we shall see that Boltzmann emphasizes even more strongly the close relations between the Second Law and probability theory. Even so, it is not always clear what these relations are exactly. Further, one may question whether his considerations of the probability of the initial state hit the nail on the head. Probability theory is equally neutral to the direction of time as is mechanics.

The true source of the reversibility problem was only identified by [Burbury, 1894a] and [Bryan, 1894] after Boltzmann's lecture in Oxford, which created an intense debate in the columns of *Nature*. They pointed out that the *Stoßzahlansatz* already contained a time-asymmetric assumption.

Indeed, this assumption requires that the number of collisions of the kind $(\vec{v}_1, \vec{v}_2) \longrightarrow (\vec{v}_1', \vec{v}_2')$ is proportional to the product $f(\vec{v}_1)f(\vec{v}_2)$ where, \vec{v}_1, \vec{v}_2 are the velocities *before* the collisions. If we would replace this by the requirement that the number of collisions is proportional to the product for the velocities \vec{v}_1', \vec{v}_2' *after* the collision, we would obtain, by a similar reasoning, $dH/dt \geq 0$. The question is now, of course, why we should prefer one assumption above the other, without falling into some kind of double standard. (I refer to [Price, 1996] for a detailed discussion of this danger.) One thing is certain, and that is that any such preference cannot be obtained from mechanics and probability theory alone.

4.4 Boltzmann (1877b): the combinatorial argument

Boltzmann's next paper (1877b) is often seen as a major departure from the conceptual basis employed in his previous work. Indeed, the conceptual shifts already indicated implicitly in his reply to Loschmidt become in this article explicit. Indeed, according to [ter Haar, 1955, p. 296] and [Klein, 1973, p. 83], it is this paper that marks the transition from kinetic theory to statistical mechanics. Further, the paper presents the famous link between entropy and 'probability' that later became known as "Boltzmann's principle", and was engraved on his tombstone as " $S = k \log W$ ".

Boltzmann begins the paper by stating that his goal is to elucidate the relationship between the Second Law and probability calculus. He notes he has repeatedly emphasized that the Second Law is related to probability calculus. In particular he points out that the 1872 paper confirmed this relationship by showing that a certain quantity [i.e. H] can only decrease, and must therefore obtain its minimum value in the state of thermal equilibrium. Yet, this connection of the Second Law with probability theory became even more apparent in his previous paper (1877a). Boltzmann states that he will now solve the problem mentioned in that paper, of calculating the probabilities of various distributions of state by determining the ratio of their numbers.

He also announces that, when a system starts in an improbable state, it will always evolve towards more probable states, until it reaches the most probable state, i.e. that of thermal equilibrium. When this is applied to the Second Law, he says, "we can identify that quantity which is usually called entropy, with the

probability of the state in question.” And: “According to the present interpretation, [the Second Law] states nothing else but that the probability of the total state of a composite system always increases” [Abh. II, pp. 165-6]. Exactly how all this is meant, he says, will become clear later in the article.

The combinatorial argument

Succinctly, and rephrased in the Ehrenfests’ terminology, the argument is as follows. Apart from Γ , the mechanical phase space containing the possible states x for the total gas system, we consider the so-called μ -space, i.e. the state space of a single molecule. For monatomic gases, this space is just a six-dimensional Euclidean space with (\vec{r}, \vec{v}) as coordinates. With each mechanical state x we can associate a collection of N points in μ -space; one for each molecule.

Now, partition μ -space into m disjoint cells: $\mu = \omega_1 \cup \dots \cup \omega_m$. These cells are taken to be rectangular in the coordinates and of equal size. Further, it is assumed that the energy of each molecule in cell ω_i in has a value ϵ_i , depending only on i . For each x , henceforth also called the *microstate* (Boltzmann’s term was the *Komplexion*), we define the *macrostate* or ‘distribution of state’ as $Z := (n_1, \dots, n_m)$, with n_i the number of particles whose molecular state is in cell ω_i . The relation between macro- and microstate is obviously non-unique since many different microstates, e.g. obtained by permuting the molecules, lead to the same macrostate. One may associate with every given macrostate Z_0 the corresponding set of microstates:

$$(56) \quad \Gamma_{Z_0} := \{x \in \Gamma : Z(x) = Z_0\}.$$

The phase space volume $|\Gamma_{Z_0}|$ of this set is proportional to the number of permutations of the particles that do not change the macrostate Z_0 . Indeed, when the six-dimensional volume of the cells ω_i is $\delta\omega$, i.e., the same for each cell, the phase space volume of the set Γ_Z is

$$(57) \quad |\Gamma_Z| = \frac{N!}{n_1! \dots n_m!} (\delta\omega)^N.$$

Moreover, assuming that $n_i \gg 1$ for all i and using the Stirling approximation for the factorials, one finds

$$(58) \quad \ln \Gamma_Z \approx N \ln N - \sum_i n_i \ln n_i + N \ln \delta\omega.$$

This expression is in fact proportional to a discrete approximation of the H -function. Indeed, putting

$$(59) \quad n_i = N f(\vec{r}_i, \vec{v}_i) \delta\omega$$

where (\vec{r}_i, \vec{v}_i) are the coordinates of a representative point in ω_i , we find

$$\sum_i n_i \ln n_i = \sum_i N f(\vec{r}_i, \vec{v}_i) \ln \left(N f(\vec{r}_i, \vec{v}_i) \delta\omega \right) \delta\omega$$

$$\begin{aligned}
 &\approx N \int f(\vec{r}, \vec{v}) \left(\ln f(\vec{r}, \vec{v}) + \ln N + \ln \delta\omega \right) d^3\vec{r} d^3\vec{v} \\
 (60) \quad &= NH + N \ln N + N \ln \delta\omega;
 \end{aligned}$$

and therefore, in view of (58):

$$(61) \quad -NH \approx \ln |\Gamma_Z|.$$

And since Boltzmann had already identified $-kNH$ with the entropy of a macrostate, one can also take entropy as proportional to the logarithm of the volume of the corresponding region in phase space. Today, $\ln |\Gamma_Z|$ is often called the *Boltzmann entropy*.

Boltzmann next considers the question for which choice of Z does the region Γ_Z have maximal size, under the constraints of a given total number of particles N , and a total energy E :

$$(62) \quad N = \sum_{i=1}^m n_i, \quad E = \sum_{i=1}^m n_i \epsilon_i.$$

This problem can easily be solved with the Lagrange multiplier technique. Under the Stirling approximation (58) one finds

$$(63) \quad n_i = \mu e^{\lambda \epsilon_i},$$

which is a discrete version of the Maxwell distribution. (Here, μ and λ are determined in terms of N and E by the constraints (62).)

Boltzmann proposes to take the macrostate with the largest volume as representing equilibrium. More generally, he also refers to these volumes as the “probability” or “permutability” of the macrostate. He therefore now expresses the Second Law as a tendency for the system to evolve towards ever more probable macrostates, until, in equilibrium, it has reached the most probable state.

Remarks and problems

1. The role of dynamics. In the present argument, no dynamical assumption has been made. In particular, it is not relevant to the argument whether the ergodic hypothesis holds, or how the particles collide. At first sight, it might seem that this makes the present argument more general than the previous one. Indeed, Boltzmann suggests at the end of the paper [Boltzmann, 1909, II p. 223] that the same argument might be applicable also to dense gases and even to solids.

However, it should be noticed that the assumption that the total energy can be expressed in the form $E = \sum_i n_i \epsilon_i$ where the energy of each particle depends only on the cell in which it is located, and not on the state of other particles is very strong. This can only be maintained, independently of the number N , if there is no interaction at all between the particles. The validity of the argument is thus really restricted to ideal gases (cf. [Uhlenbeck and Ford, 1963]).

2. The choice of cells. One might perhaps hope, at first sight, that the procedure of partitioning μ -space into cells is only a technical or didactic device and can be eliminated by finally taking a limit in which $\delta\omega \rightarrow 0$; similar to the procedure of his 1872 paper. This hope is dashed because the expression (58) diverges. Indeed, the whole prospect of using combinatorics would disappear if we did not adopt a finite partition. But also the special choice to give all cells equal volume in position and velocity variables is not quite self-evident, as Boltzmann himself shows. In fact, before he develops the argument given here, his paper presents a discussion in which the particles are characterized by their energy instead of position and velocity. This leads him to carve up μ -space into cells of equal size $\delta\epsilon$ in energy. He then shows that the combinatorial argument *fails* to reproduce the desired Maxwell distribution for particles moving in 3 spatial dimensions.³⁵ This failure is then remedied [Boltzmann, 1909, II, p. 190] by switching to a choice of equally sized cells in $\delta\omega$ in position and velocity. The latter choice is apparently 'right', in the sense that leads to the desired result. However, since the choice clearly cannot be relegated to a matter of convention, it leaves open the question of justification.

Modern commentators are utterly divided in the search for a direction in which a motivation for the choice of the size of these cells can be found. Some argue that the choice should be made in accordance with the actual finite resolution of measuring instruments or human observation capabilities. The question whether these do in fact favour a partition into cells of equal phase space volume has hardly been touched upon. Others [Popper, 1982; Redhead, 1995] reject an appeal to observation capacities on the grounds that these would introduce a 'subjective' or 'anthropocentric' element into the explanation of irreversibility (see also [Jaynes, 1965; Grünbaum, 1973; Denbigh and Denbigh, 1985; Ridderbos, 2002]).

3. Micro versus macro. The essential step in the argument is the distinction between micro- and macrostates. This is indeed the decisive new element, that allowed Boltzmann a complete reinterpretation of the notion and role of probability.

In 1872 and before, the distribution of state f was *identified* with a probability (namely of a molecular state, cf. Remark 1 of paragraph 4.2). On the other hand, in the present work it, or its discrete analogue Z , is a description of the macrostate of the gas, to which a probability is *assigned*. Essentially, the role of the distribution of state has been shifted from defining a probability measure to being a stochastic variable. Its previous role is taken over by a new idea: Probabilities are not assigned to the particles, but to the macrostate of the gas as a whole, and measured by the corresponding volume in phase space.

Another novelty is that Boltzmann has changed his concept of equilibrium. Whereas previously the defining characteristic of equilibrium was its stationarity,

³⁵The problem is that for an ideal gas, where all energy is kinetic, $\delta\epsilon \propto v\delta v$. On the other hand, for three-dimensional particles, $\delta\omega \propto v^2\delta v$. The function f derived from (59) and (63) thus has a different dependence on v in the two cases. As Boltzmann notes, the two choices are compatible for particles in two dimensions (i.e. discs moving in a plane).

in Boltzmann's new view it is conceived as the macrostate (i.e. a region in phase space) that takes up the largest volume. As a result, a system in a state of equilibrium need not remain there: in the course of time, the microstate of the system may fluctuate in and out of this equilibrium region. Boltzmann briefly investigated the probability of such fluctuations in his [Boltzmann, 1878]. Almost thirty years later, the experimental predictions for fluctuation phenomena by Einstein and Smoluchowski provided striking empirical successes for statistical mechanics.

4. But what about evolutions? Perhaps the most important issue is this. What exactly is the relation of the 1877b paper to Loschmidt's objection and Boltzmann's primary reply to it (1877a)? The primary reply (cf. paragraph 4.3) can be read as an announcement of two subjects of further investigation:

From the relative numbers of the various distributions of state, one might even be able to calculate their probabilities. This could lead to an interesting method of determining thermal equilibrium [Boltzmann, 1909, II, p. 121]

This is a problem about equilibrium. The second announcement was that Boltzmann said "The case is completely analogous for the Second Law" [Boltzmann, 1909, II, p. 121]. Because there are so very many more uniform than non-uniform distributions, it should be extraordinarily improbable that a system should evolve from a uniform distribution of states to a non-uniform distribution of states. This is a problem about evolution (cf. point 3 of section 4.3). In other words, one would like to see that something like the statistical H -theorem actually holds.

Boltzmann's [1877b] is widely read as a follow-up to these announcements. Indeed, Boltzmann repeats the first quote above in the introduction of the paper [Boltzmann, 1909, II, p. 165], indicating that he will address this problem. And so he does, extensively. Yet he also states:

Our main goal is not to linger on a discussion of thermal equilibrium, but to investigate the relations of probability with the Second Law of thermodynamics [Boltzmann, 1909, II, p. 166].

Thus, the main goal of [1877b] is apparently to address the problem concerning evolutions and to show how they relate to the Second Law. Indeed, this is what one would naturally expect since the reversibility objection is, after all, a problem concerned with evolutions. Even so, a remarkable fact is that the 1877b paper hardly ever touches its self-professed "main goal" at all. As a matter of fact, I can find only one passage in the remainder of the paper where a connection with the Second Law is mentioned.

It occurs in Section V [Boltzmann, 1909, II, p. 216-7]. After showing that in equilibrium states for monatomic gases the 'permutability measure' $\ln |\Gamma_Z|$ (for which Boltzmann's notation is Ω) is proportional to the thermodynamical entropy,

up to an arbitrary additive constant, he concludes that, by choosing the constant appropriately:³⁶

$$(64) \quad \int \frac{dQ}{T} = \frac{2}{3} \Omega \left[= \frac{2}{3} \ln |\Gamma_Z| \right]$$

and adds:

It is known that when a system of bodies goes through reversible changes, the total sum of the entropies of all these bodies remains constant; but when there are among these processes also irreversible (nicht umkehrbar) changes, then the total entropy must necessarily increase. This follows from the familiar circumstance that $\int dQ/T$ is negative for an irreversible cyclic process. In view of (64), the sum of all permutability measures of all bodies $\sum \Omega$, or their total permutability measure, must also increase. Hence, permutability is a quantity which is, up to a multiplicative and additive constant, identical to entropy, but which retains a meaning also during the passage of an irreversible body [sic— read: “process”], in the course of which it continually increases [Boltzmann, 1909, II p.217]

How does this settle the problem about evolutions, and does it provide a satisfactory refutation of the reversibility objection? In the literature, there are at least four views about what Boltzmann’s response actually intended or accomplished.

4 α . Relying on the separation between micro- and macroscales: A view that has been voiced recently, e.g. by [Goldstein, 2001], is that Boltzmann had, by his own argument, adequately and straightforwardly explained why entropy should tend to increase. In particular, this view argues, the fact of the overwhelmingly large phase space volume of the set Γ_{eq} of all equilibrium phase points, compared to the set of non-equilibrium points already provides a sufficient argument.

For a non-equilibrium phase point x of energy E , the Hamiltonian dynamics governing the motion x_t arising from x would have to be ridiculously special to avoid reasonably quickly carrying x_t into Γ_{eq} and keeping it there for an extremely long time — unless, of course x itself were ridiculously special [Goldstein, 2001, p. 6].

In fact, this view may lay some claim to being historically faithful. As we have seen, [Boltzmann, 1877a] did claim that the large probability for an evolution towards equilibrium did follow from the large differences in number of states.

The main difficulty with this view is that, from a modern perspective, it is hard to maintain that it is adequate. States don’t evolve into other states just because there are more of the latter, or because they make up a set of larger

³⁶Actually, equation (64) is the closest he got to the famous formula on his tombstone, since $\Omega = \ln W$, and Boltzmann adopts a temperature scale that makes $k = 2/3$.

measure. The evolution of a system depends only on its initial state and its Hamiltonian. Questions about evolution can only be answered by means of an appeal to dynamics, not by the measure of sets alone. To take an extreme example, the trajectory covered by x_t , i.e. the set $\{x_t : t \in \mathbb{R}\}$ is a set of measure zero anyway; and hence very special. By contrast, its complement, i.e. the set of states *not* visited by a given trajectory is huge: it has measure one. Certainly, we cannot argue that the system cannot avoid wandering into the set of states that it does not visit. Another example is that of a system of non-interacting particles, e.g., the ideal gas. In this case, all the energies of the individual particles are conserved, and because of these conserved quantities, the phase point can only visit a very restricted region of phase space.³⁷

The lesson is, of course, that in order to obtain any satisfactory argument why the system should tend to evolve from non-equilibrium states to the equilibrium state, we should make some assumptions about its dynamics. In any case, judgments like “reasonable” or “ridiculous” remain partly a matter of taste. The reversibility objection is a request for mathematical proof (which, as the saying goes, is something that even convinces an unreasonable person).

4β. Relying on the ergodic hypothesis: A second, and perhaps the most well-known, view to this problem is the one supplied by the Ehrenfests. In essence, they suggest that Boltzmann somehow relied on the ergodic hypothesis in his argument.

It is indeed evident that if the ergodic hypothesis holds, a state will spend time in the various regions of the energy hypersurface in phase space in proportion to their volume. That is to say, during the evolution of the system along its trajectory, regions with a small volume, corresponding to highly non-uniform distributions of state are visited only sporadically, and regions with larger volume, corresponding to more uniform distributions of state more often.

This should also make it plausible that if a system starts out from a very small region (an improbable state) it will display a tendency to evolve towards the overwhelmingly larger equilibrium state. Of course, this ‘tendency’ would have to be interpreted in a qualified sense: the same ergodic hypothesis would imply that the system cannot stay inside the equilibrium state forever and thus there would necessarily be fluctuations in and out of equilibrium. Indeed, one would have to state that the tendency to evolve from improbable to probable states is itself a probabilistic affair: as something that holds true for most of the initial states, or for most of the time, or as some or other form of average behaviour. In short, we would then hopefully obtain some statistical version of the H -theorem. What exactly the statistical H -theorem should say remains an open problem in the Ehrenfests’ point of view. Indeed they distinguish between several interpretations (the so-called ‘concentration curve’ and the ‘bundle of H -curves’ [Ehrenfest and Ehrenfest-Afanassjewa, 1912, p. 31–35]).

³⁷It is somewhat ironic to note, in view of remark 1 above, that this is the only case compatible with Boltzmann’s argument. This gives rise to Khinchin’s “methodological paradox” (cf. 1019).

Now, it is undeniable that the Ehrenfests' reading of Boltzmann's intentions has some clear advantages. In particular, even though nobody has yet succeeded in proving a statistical H -theorem on the basis of the ergodic hypothesis, or on the basis of the assumption of metric transitivity (cf. paragraph 6.1, one might hope that some statistical version of the H -theorem is true.

One problem here is that the assumptions Boltzmann used in his paper are restricted to non-interacting molecules, for which the ergodic hypothesis is demonstrably false. But even more importantly, it is clear that Boltzmann did not follow this line of argument in 1877b at all. Indeed, he nowhere mentions the ergodic hypothesis. In fact he later commented on the relation between the 1877b paper and the ergodic hypothesis of 1868, saying:

On that occasion [i.e. in (1877b)] . . . I did not wish to touch upon the question whether a system is capable of traversing all possible states compatible with the equation of energy [Boltzmann, 1881a, Abh. II p. 572].

4γ. Relying on the H -theorem: A third point of view, one to which this author adhered until recently, is that, in (1877b) Boltzmann simply relied on the validity of the H -theorem of 1872. After all, it was the 1872 paper that proposed to interpret $-NH$ as entropy (modulo multiplicative and additive constants), on the basis of the alleged theorem that it could never decrease. The 1877b paper presents a new proposal, to link the entropy of a macrostate with $\ln |\Gamma_Z|$. But this proposal is motivated, if not derived, by showing that $\ln |\Gamma_Z|$ is (approximately) equal to $-NH$, as in (61), whose interpretation as entropy was established in (1872). It thus seems plausible to conjecture that Boltzmann's thinking relied on the results of that paper, and that the claim that states will evolve from improbable to probable states, i.e. that $\ln |\Gamma_Z|$ shows a tendency to increase in time, likewise relied on the H -theorem he had proved there.³⁸ The drawback of this reading is that it makes Boltzmann's response to the reversibility objection quite untenable. Since the objection as formulated in his (1877a) calls the validity of the H -theorem into question, a response that *presupposes* the validity of this theorem is of no help at all.

4δ. Bypassing the H -theorem: [Janssen, 2002] has a different reading. He notes: "In Boltzmann's 1877 paper the statement that systems never evolve from more probable to less probable states is presented only as a new way of phrasing the Second Law, not as a consequence of the H -theorem" (p. 13). Indeed, any explicit reference to the H -theorem is absent in the 1877b paper. However, what we are to make of this is not quite certain. The earlier paper (1877a) did not mention the theorem either, but only discussed "any attempt to prove that $\int \frac{dQ}{T} \leq 0$ ". Still, this is commonly seen as an implicit reference to what is now known as the

³⁸The conjecture is supported by the fact Boltzmann's later exposition in 1896 is presented along this line.

H -theorem, but which did not yet have a particular name at that time. Indeed, the H -theorem itself was characterized in 1872 only as a new proof that $\int \frac{dQ}{T} \leq 0$ (cf. the quotation on page 965). So, the fact that the H -theorem is not explicitly mentioned in (1877b) is not by itself a decisive argument that he did not intend to refer to it.

Even so, the fact that he presented the increase of entropy as something which was well-known and did not refer to the 1872 paper at all, does make Janssen's reading plausible. So, perhaps Boltzmann merely relied on the empirical validity of the Second Law as a ground for this statement, and not at all on any proposition from kinetic theory of gases.³⁹ This, of course, would undermine even more strongly the point of view that Boltzmann had a statistical version of the H -theorem, or indeed any theorem at all, about the probability of time evolution.

The reversibility objection was not about a relationship between the phenomenological Second Law and the H -theorem, but about the relationship between the H -theorem and the mechanical equations of motion. So even though Janssen's reading makes Boltzmann's views consistent, it does not make the 1877b paper provide a valid answer to Loschmidt's objection.

4ε. The urn analogy — victory by definition? At the risk of perhaps overworking the issue, I also want to suggest a fifth reading. Boltzmann's (1877b) contains an elaborate discussion of repeated drawings from an urn. In modern terms, he considers a Bernoulli process, i.e., a sequence of independent identically distributed repetitions of an experiment with a finite number of possible outcomes. To be concrete, consider an urn filled with m differently labeled lots, and a sequence of N drawings, in which the lot i is drawn n_i times ($\sum_{i=1}^m n_i = N$). He represents this sequence by a "distribution of state" $Z = (n_1, \dots, n_m)$. In this discussion, the probability of these distributions of state is at first identified with the (normalized) number of permutations by which Z can be produced. In other words

$$(65) \text{ Prob}(Z) \propto \frac{N!}{n_1! \cdots n_m!}.$$

But halfway this discussion [Boltzmann, 1909, II, p. 171], he argues that one can redefine probabilities in an alternative fashion, namely, as the relative frequency of occurrence during *later* drawings of a sequence of N lots. Thus, even when, on a particular trial, an improbable state Z occurred, we can still argue that on a later drawings, a more probable state will occur. Boltzmann speaks about the changes in Z during the consecutive repetitions as an *evolution*. He then says:

³⁹Further support for this reading can be gathered from later passages. For example, [Boltzmann, 1897b] writes "Experience shows that a system of interacting bodies is always found 'initially' in an improbable state and will soon reach the most probable state (that of equilibrium). [Boltzmann, 1909, III, p. 607]. Here too, Boltzmann presents the tendency to evolve from improbable to more probable states as a fact of experience rather than the consequence of any theorem.

The most probable distribution of state must therefore be defined as that one to which most [states] will evolve to [Boltzmann, 1909, II, p. 172].

Although he does not make the point quite explicitly, the discussion of urn drawings is undoubtedly meant as an analogy for the evolution of the distribution of state in a gas. Hence, it is not implausible that, in the latter case too, Boltzmann might have thought that *by definition* the most probable distribution of state is the one that most states will evolve to. And this, in turn, would mean that he regarded the problem about evolutions not as something to be proved, and that might depend on the validity of specific dynamical assumptions like the ergodic hypothesis or the *Stoßzahlansatz*, but as something already settled from the outset. This would certainly explain why Boltzmann did not bother to address the issue further.

Even so, this reading too has serious objections. Apart from the fact that it is not a wise idea to redefine concepts in the middle of an argument, the analogy between the evolution of an isolated gas and a Bernoulli process is shaky. In the first case, the evolution is governed by deterministic laws of motion; in the latter one simply avoids any reference to underlying dynamics by the stipulation of the probabilistic independence of repeated drawings. However, see paragraph 6.2.

To sum up this discussion of Boltzmann's answer to the reversibility objection: it seems that on all above readings of his two 1877 papers, the lacuna between what Boltzmann had achieved and what he needed to do to answer Loschmidt satisfactorily — i.e. to address the issue of the evolution of distributions of state and to prove that non-uniform distributions tend, in some *statistical* sense, to uniform ones, or to prove any other reformulation of the *H*-theorem — remains striking.

4.5 *The recurrence objection*

Poincaré

In 1890, in his famous treatise on the three-body problem of celestial mechanics, Poincaré derived what is nowadays called the recurrence theorem. Roughly speaking, the theorem says that for every mechanical system with a bounded phase space, almost every initial state of the system will, after some finite time, return to a state arbitrarily closely to this initial state, and indeed repeat this infinitely often.

In modern terms, the theorem can be formulated as follows:

RECURRENCE THEOREM: Consider a dynamical system⁴⁰ $\langle \Gamma, \mathcal{A}, \mu, T \rangle$

⁴⁰See section 6.1 for a definition of dynamical systems. But in short: Γ is a phase space, \mathcal{A} a family of measurable subsets of Γ and T is a one-parameter continuous group of time evolutions $T_t : \Gamma \times \mathbb{R} \rightarrow \Gamma$.

with $\mu(\Gamma) < \infty$. Let $A \in \mathcal{A}$ be any measurable subset of Γ , and define, for a given time τ , the set

$$(66) \quad B = \{x : x \in A \ \& \ \forall t \geq \tau : T_t x \notin A\}$$

Then

$$(67) \quad \mu(B) = 0.$$

In particular, for a Hamiltonian system, if we choose Γ to be the energy hypersurface Γ_E , take A to be a ‘tiny’ region in Γ_E , say an open ball of diameter ϵ in canonical coordinates, the theorem says that the set of points in this region whose evolution is such that they will, after some time τ , never return to region A , has measure zero. In other words, almost every trajectory starting within A will after any finite time we choose, later return to A .

Poincaré had already expressed his objections against the tenability of a mechanical explanation of irreversible phenomena in thermodynamics earlier (e.g. [Poincaré, 1889]). But armed with his new theorem, he could make the point even stronger. In his 1893, he argued that the mechanical conception of heat is in contradiction with our experience of irreversible processes. According to the English kinetic theories, says Poincaré:

[t]he world tends at first towards a state where it remains for a long time without apparent change; and this is consistent with experience; but it does not remain that way forever, if the theorem cited above is not violated; it merely stays there for an enormously long time, a time which is longer the more numerous are the molecules. This state will not be the final death of the universe but a sort of slumber, from which it will awake after millions and millions of centuries.

According to this theory, to see heat pass from a cold body into a warm one, it will not be necessary to have the acute vision, the intelligence and the dexterity of Maxwell’s demon; it will suffice to have a little patience [Brush, 2003, p.380].

He concludes that these consequences contradict experience and lead to a “definite condemnation of mechanism” [Brush, 2003, p.381].

Of course, Poincaré’s “little patience”, even for “millions and millions of centuries” is a rather optimistic understatement. Boltzmann later estimated the time needed for a recurrence in 1 cc of air to be $10^{10^{19}}$ seconds (see below): utterly beyond the bounds of experience. Poincaré’s claim that the results of kinetic theory are contradicted by experience is thus too hasty.

Poincaré’s article does not seem to have been noticed in the contemporary German-language physics community — perhaps because he criticized English theories only. However, Boltzmann was alerted to the problem when a slightly different argument was put forward by Zermelo in 1896. The foremost difference is that in Zermelo’s argument experience does not play a role.

Zermelo's argument

Zermelo (1896a) points out that for a Hamiltonian mechanical system with a bounded phase space, Poincaré's theorem implies that, apart from a set of singular states, every state must recur almost exactly to its initial state, and indeed repeat this recurrence arbitrarily often. As a consequence, for any continuous function F on phase space, $F(x_t)$ cannot be monotonically increasing in time, (except when the initial state is singular); whenever there is a finite increase, there must also be a corresponding decrease when the initial state recurs. (see [Olsen, 1993] for a modern proof of this claim) Thus, it would be impossible to obtain 'irreversible' processes. Along the way, Zermelo points out a number of options to avoid the problem.

1. Either we assume that the gas system has no bounded phase space. This could be achieved by letting the particles reach infinite distances or infinite velocities. The first option is however excluded by the assumption that a gas is contained in a finite volume. The second option could be achieved when the gas consists point particles which attract each other at small distances, (e.g. an $F \propto r^{-2}$ inter-particle attractive force can accelerate them toward arbitrarily high velocities.) However, on physical grounds one ought to assume that there is always repulsion between particles at very small distances.

2. Another possibility is to assume that the particles act upon each other by velocity-dependent forces. This, however would lead either to a violation of the conservation of energy or the law of action and reaction, both of which are essential to atomic theory.

3. The H -theorem holds only for those special initial states which are the exception to the recurrence theorem, and we assume that only those states are realized in nature. This option would be unrefutable, says Zermelo. Indeed, the reversibility objection has already shown that not all initial states can correspond to the Second Law. However, here we would have to exclude the overwhelming majority of all imaginable initial states, since the exceptions to the Recurrence Theorem only make up a set of total extension (i.e. in modern language: measure) zero. Moreover, the smallest change in the state variables would transform a singular state into a recurring state, and thus suffice to destroy the assumption. Therefore, this assumption "would be quite unique in physics and I do not believe that anyone would be satisfied with it for very long."

This leaves only two major options:

4. The Carnot-Clausius principle must be altered.⁴¹
5. The kinetic theory must be formulated in an essentially different way, or even be given up altogether.

Zermelo does not express any preference between these last two options. He concludes that his aim has been to explain as clearly as possible what can be proved rigorously, and hopes that this will contribute to a renewed discussion and

⁴¹By this term, Zermelo obviously referred to the Second Law, presumably including the Zeroth Law.

final solution of the problem.

I would like to emphasize that, in my opinion, Zermelo's argument is entirely correct. If he can be faulted for anything, it is only that he had not noticed that in his very recent papers, Boltzmann had already been putting a different gloss on the H -theorem.

Boltzmann's response

[Boltzmann, 1896b] response opens by stating that he had repeatedly pointed out that the theorems of gas are statistical. In particular, he says, he had often emphasized as clearly as possible that the Maxwell distribution law is not a theorem from ordinary mechanics and cannot be proven from mechanical assumptions.⁴² Similarly, from the molecular viewpoint, the Second Law appears merely as a probability statement. He continues with a sarcastic remark:

Zermelo's paper shows that my writings have been misunderstood; nevertheless it pleases me for it appears to be the first indication that these works have been noticed in Germany.⁴³

Boltzmann agrees that Poincaré's recurrence theorem is "obviously correct", but claims that Zermelo's application of the theorem to gas theory is not. His counter argument is very similar to his (1895) presentation in *Nature*, a paper that Zermelo had clearly missed.

In more detail, this argument runs as follows. Consider a gas in a vessel with perfectly smooth and elastic walls, in an arbitrary initial state and let it evolve in the course of time. At each time t we can calculate $H(t)$. Further, consider a graph of this function, which Boltzmann called: *the H-curve*. In his second reply to Zermelo [Boltzmann, 1897a], he actually produced a diagram. A rough and modernized version of such an H -curve is sketched in Fig. 3.

Barring all cases in which the motion is 'regular', e.g. when all the molecules move in one common plane, Boltzmann claims the following properties of the curve:

- (i). For most of the time, $H(t)$ will be very close to its minimum value, say H_{\min} . Moreover, whenever the value of $H(t)$ is very close to H_{\min} , the distribution of molecular velocities deviates only very little from the Maxwell distribution.

⁴²This is, as we have seen, a point Boltzmann had been making since 1877. However, one might note that just a few years earlier, [Boltzmann, 1892], after giving yet another derivation of the Maxwell distribution (this time generalized to a gas of hard bodies with an arbitrary number of degrees of freedom that contribute quadratic terms to the Hamiltonian), had concluded: "I believe therefore that its correctness [i.e. of the Maxwell distribution law] as a theorem of analytical mechanics can hardly be doubted" [Boltzmann, 1909, III p.432]. But as we have seen on other occasions, for Boltzmann, statements that some result depended essentially on probability theory, and the statement that it could be derived as a mechanical theorem, need not exclude each other.

⁴³Eight years earlier, Boltzmann had been offered the prestigious chair in Berlin as successor of Kirchhoff, and membership of the Prussian Academy. The complaint that his works did not draw attention in Germany is thus hard to take seriously.




Figure 3. A (stylized) example of an H -curve

- (ii). The curve will occasionally, but very rarely, rise to a peak or summit, that may be well above H_{\min} .
- (iii). The probability of a peak decreases extremely rapidly with its height.

Now suppose that, at some initial time $t = 0$, the function takes a very high value H_0 , well above the minimum value. Then, Boltzmann says, it will be enormously probable that the state will, in the course of time, approach the Maxwell distribution, i.e., $H(t)$ will decrease towards H_{\min} ; and subsequently remain there for an enormously long time, so that the state will deviate only very little from the Maxwell distribution during vanishingly short durations. Nevertheless, if one waits even longer, one will encounter a new peak, and indeed, the original state will eventually recur. In a mathematical sense, therefore, these evolutions are periodic, in full conformity with Poincaré's recurrence theorem.

What, then, is the failure of Zermelo's argument? Zermelo had claimed that only very special states have the property of continually approaching the Maxwell distribution, and that these special states taken together make up an infinitely small number compared to the totality of possible states. This is incorrect, Boltzmann says. For the overwhelming majority of states, the H -curve has the qualitative character sketched above.

Boltzmann also took issue with (what he claimed to be Zermelo's) conclusion that the mechanical viewpoint must somehow be changed or given up. This conclusion would only be justified, he argues, if this viewpoint led to some consequence that contradicted experience. But, Boltzmann claims, the duration of the recurrence times is so large that no one will live to observe them.

To substantiate this claim about the length of the recurrence time, he presents, in an appendix an estimate of the recurrence time for 1 cc of air at room tempera-

ture and pressure. Assuming there are 10^9 molecules in this sample,⁴⁴ and choosing cells in the corresponding μ -space as six-dimensional cubes of width 10^{-9} m in (physical) space and 1 m/s in velocity space, Boltzmann calculates the number of different macrostates, i.e. the number of different ways in which the molecules can be distributed over these cells as (roughly) 10^{10^9} . He then assumes that, before a recurrence of a previous macrostate, the system has to pass through *all* other macrostates. Even if the molecules collide very often, so that the system changes its macrostate 10^{27} times per second, the total time it takes to go through this huge number of macrostates will still take $10^{10^9-27} \approx 10^{10^9}$ seconds. In fact, this time is so immensely large that its order of magnitude is not affected whether we express it in seconds, years, millennia, or what have you.

The upshot is, according to Boltzmann: if we adopt the view that heat is a form of motion of the molecules, obeying the general laws of mechanics, and assume that the initial state of a system is very unlikely, we arrive at a theorem which corresponds to the Second Law for all observed phenomena. He ends with another sarcasm:

All the objections raised against the mechanical view of Nature are therefore empty and rest on errors. But whoever cannot overcome the difficulties, which a clear understanding of the theorems of gas theory poses, should indeed follow the advice of Mr Zermelo and decide to give up the theory completely. [Boltzmann, 1909, III p. 576].

Zermelo's reply

[Zermelo, 1896b] notes that Boltzmann's response confirms his views by admitting that the Poincaré theorem is correct and applicable to a closed system of gas molecules. Hence, in such a system, "all [sic] motions are *periodic* and not *irreversible* in the strict sense". Thus, kinetic gas theory cannot assert that there is a strict monotonic increase of entropy as the Second Law would require. He adds: "I think this general clarification was not at all superfluous" [Brush, 2003, p. 404].

Therefore, Zermelo argues, his main point had been conceded: there is indeed a conflict between thermodynamics and kinetic theory, and it remains a matter of taste which of the two is abandoned. Zermelo admits that observation of the Poincaré recurrences may well fall beyond the bounds of human experience. He points out (correctly) that Boltzmann's estimate of the recurrence time presupposes that the system visits *all* other cells in phase space before recurring to an initial state. This estimate is inconclusive, since the latter assumption is somewhat ad hoc. In general, these recurrence times need not "come out so 'comfortingly' large" [Brush, 2003, p. 405]. But, as I stressed before, the relation with experience simply was no issue in Zermelo's objection.

⁴⁴Actually, modern estimates put the number of molecules in 1cc of air closer to 10^{19} , which would make Boltzmann's estimate for recurrence time even larger still, i.e. $10^{10^{19}}$.

The main body of Zermelo's reply is taken by an analysis of the justification of and consequences drawn from Boltzmann's assumption that the initial state is very improbable, i.e., that H_0 is very high. Zermelo argues that even in order to obtain an approximate or empirical analogue of the Second Law, as Boltzmann envisaged, i.e. an approach to a long-lasting, but not permanent equilibrium state, it would not suffice to show this result for one particular initial state. Rather, one would have to show that evolutions *always* take place in the same sense, at least during observable time spans.

As Zermelo understands it, Boltzmann does not merely assume that the initial state has a very high value for H , but also that, as a rule, the initial state lies on a maximum, or has just passed a maximum. If this assumption is granted, then it is obvious that one can only observe a decreasing flank of the H -curve. However, Zermelo protests, one could have chosen any time as the initial time. In order to obtain a satisfactorily general result, the additional assumption would thus have to apply at all times. But then the H -curve would have to consist entirely of maxima. But this leads to nonsense, Zermelo argues, since the curve cannot be constant. Zermelo concludes that Boltzmann's assumptions about the initial state are thus in need of further *physical* explanation.

Further, Zermelo points out that probability theory, by itself, is neutral with respect to the direction of time, so that no preference for evolutions in a particular sense can be derived from it. He also points out that Boltzmann apparently equates the duration of a state and its extension (i.e. the relative time spent in a region and the relative volume of that region in phase space). "I cannot find that he has actually *proved* this property" [Brush, 2003, p. 406].

Boltzmann's second reply

In his second reply 1897a, Boltzmann rebuts Zermelo's demand for a physical explanation of his assumptions about the initial state of the system with the claim that the question is not what will happen to an arbitrarily chosen initial state, but rather what will happen to a system in the present state of the universe.

He argues that one should depart from the (admittedly unprovable) assumption that the universe (or at least a very large part of the universe that surrounds us) started in a very improbable state and still is in an improbable state. If one then considers a small system (e.g. a gas) that is suddenly isolated from the rest of the universe, there are the following possibilities: (i) The system may already be in equilibrium, i.e. H is close to its minimum value. This, Boltzmann says, is by far the most probable case. But among the few cases in which the system is not in equilibrium, the most probable case is (ii) that H will be on a maximum of the H -curve, so that it will decrease in both directions of time. Even more rare is the case in which (iii) the initial value of H will fall on a decreasing flank of the H curve. But such cases are just as frequent as those in which (iv) H falls on an increasing flank.⁴⁵

⁴⁵The Ehrenfests 1912 later added a final possible case (v): H may initially be on a local

Thus, Boltzmann's explanation for the claim that H is initially on a maximum is that this would be the most likely case for a system not in equilibrium, which is isolated from the rest of the universe in its present state.

This occasion is perhaps the first time that Boltzmann advanced an explanation of his claims as being due to an assumption about initial state of the system, ultimately tied to an assumption about the initial conditions of the universe. Today, this is often called the *past-hypothesis* (cf. [Albert, 2000; Winsberg, 2004; Callender, 2004; Earman, 2006]).

He ends his reply with the observation that while the mechanical conception of gas theory agrees with the Clausius-Carnot conception [i.e. thermodynamics] in all observable phenomena, a virtue of the mechanical view is that it might eventually predict new phenomena, in particular for the motion of small bodies suspended in fluids. These prophetic words were substantiated eight years later in Einstein's work on Brownian motion.

However, he does not respond to Zermelo's requests for more definite proofs of the claims (1)–(3), or of the equality of phase space volume and time averages in particular. He bluntly states that he has thirty years of priority in measuring probabilities by means of phase space volume (which is true) and adds that he has always had done so (which is false). Even so, one cannot interpret this claim of Boltzmann as a rejection of the time average conception of probability. A few lines below, he claims that the most probable states will also occur most frequently, except for a vanishingly small number of initial states. He does not enter into a proof of this. Once again, this provides an instance where the Ehrenfests conjectured that Boltzmann might have had the ergodic hypothesis in the back of his mind.

Remarks

Boltzmann's replies to Zermelo have been recommended as "superbly clear and right on the money" [Lebowitz, 1999, p. S347]. However, as will clear from the above and the following remarks, I do not share this view. See also [Klein, 1973; Curd, 1982; Batterman, 1990; Cercignani, 1998; Brush, 1999; Earman, 2006] for other commentaries on the Zermelo-Boltzmann dispute.

1. The issues at stake It is clear that, in at least one main point of the dispute, Boltzmann and Zermelo had been talking past each other. When Zermelo argued that in the kinetic theory of gases there can be no continual approach towards a final stationary state, he obviously meant this in the sense of a limit $t \rightarrow \infty$. But Boltzmann's reply indicates that he took the "approach" as something that is not certain but only probable, and as lasting for a very long, but finite time. His graph of the H -curve makes abundantly clear that $\lim_{t \rightarrow \infty} H(t)$ does not exist.

It is true that his statistical reading of the H -theorem, as laid down in the claims (1)–(3) above, was already explicit in (Boltzmann 1895), and thus Boltzmann could

minimum of the H -curve, so that it increases in both directions of time. But by a similar reasoning, that case is even less probable than the cases mentioned by Boltzmann.

claim with some justification that his work had been overlooked. But in fairness, one must note that, even in this period, Boltzmann was sending mixed messages to his readers. Indeed, the first volume of Boltzmann's *Lectures on Gas Theory*, published in 1896, stressed, much like his original [1872] paper on the H -theorem, the necessity and exceptionless generality of the H -theorem, adding only that the theorem depended on the assumption of molecular disorder (as he then called the *Stoßzahlansatz*):⁴⁶ “

[T]he quantity designated as H can only decrease; at most it can remain constant.[...] The only assumption we have made here is that the distribution of velocities was initially ‘molecularly disordered’ and remains disordered. Under this condition we have therefore proved that the quantity called H can only decrease and that the distribution of velocities must necessarily approach the Maxwell distribution ever more closely [Boltzmann, 1896, § 5, p. 38].

Zermelo might not have been alone in presuming that Boltzmann had intended this last claim literally, and was at least equally justified in pointing out that Boltzmann's clarification “was not at all superfluous”.

On the other hand, Boltzmann misrepresented Zermelo's argument as concluding that the mechanical view should be given up. As we have seen, Zermelo only argued for a *dilemma* between the strict validity of the kinetic theory and the strict validity of thermodynamics. Empirical matters were not relevant to Zermelo's analysis. Still, Boltzmann is obviously correct when he says that the objection does not yet unearth a conflict with experience. Thus, his response would have been more successful as a counter-argument to Poincaré than to Zermelo.

2. The statistical reading of the H -theorem. Another point concerns the set of claims (1)–(3) that Boltzmann lays down for the behaviour of the H -curve. Together, they form perhaps the most clearly stated and explicit form of the “statistical reading of the H -theorem” (cf remark 3 on page 972). Yet they only have a loose connection to the original theorem. It is unclear, for example, whether these claims still depend on the *Stoßzahlansatz*, the assumption that the gas is dilute, etc. It thus remains a reasonable question what argument we have for their validity. Boltzmann offers none. In his 1895 paper in *Nature*, he argued as if he had proved as much in his earlier papers, and added tersely: “I will not here repeat the proofs given in my papers” [Boltzmann, 1909, III p. 541]. But surely, Boltzmann never proved anything concerning the probability of the time evolution of H , and at this point there remains a gap in his theory. Of course, one might

⁴⁶in his reply to Zermelo, Boltzmann claimed that his discussion of the H -theorem in the *Lectures on Gas theory* was intended under the explicitly emphasized assumption that the number of molecules was infinite, so that the recurrence theorem did not apply. However, I can find no mention of such an assumption in this context. On the contrary, the first occasion on which this latter assumption appears is in §6 on page 46 where it is introduced as “an assumption we shall make later”, suggesting that the previous discussion did *not* depend on it.

speculate on ways to bridge this gap; e.g. that Boltzmann implicitly and silently relied on the ergodic hypothesis, as the Ehrenfests suggested or in other ways, but I refrain from discussing this further. The most successful modern attempt so far to formulate and prove a statistical H -theorem has been provided by Lanford, see paragraph 6.4 below.

5 GIBBS' STATISTICAL MECHANICS

The birth of statistical mechanics in a strict sense, i.e. as a coherent and systematic theory, is marked by the appearance of J.W. Gibbs's book (1902) which carries this title: *Elementary Principles in Statistical Mechanics; developed with especial reference to the rational foundation of thermodynamics*. His point of departure is a general mechanical system governed by Hamiltonian equations of motion, whose (micro)states are represented by points in the mechanical phase space Γ .

Gibbs avoids specific hypotheses about the microscopic constitution of such a system. He refers to the well-known problem concerning the anomalous values of the specific heat for gases consisting of diatomic molecules (mentioned in footnote 10), and remarks:

Difficulties of this kind have deterred the author from attempting to explain the mysteries of nature, and have forced him to be contented with the more modest aim of deducing some of the more obvious propositions relating to the statistical branch of mechanics [Gibbs, 1902, p. viii].

It is clear from this quote that Gibbs' main concern was with logical coherence, and less with the molecular constitution. (Indeed, only the very last chapter of the book is devoted to systems composed of molecules.) This sets his approach apart from Maxwell and Boltzmann.⁴⁷

The only two ingredients in Gibbs' logical scheme are mechanics and probability. Probability is introduced here as an ingredient not reducible to the mechanical state of an individual system, but by means of the now familiar "ensemble":

We may imagine a great number of systems of the same nature, but differing in the configurations and velocities which they have at a given instant, and differing not merely infinitesimally, but it may be so as to embrace every conceivable combination of configuration and velocities. And here we may set the problem, not to follow a particular system through its succession of configurations, but to determine how the whole number of systems will be distributed among the various conceivable configurations and velocities at any required time, when the distribution has been given for some one time [Gibbs, 1902, p. v].

⁴⁷It also sets him apart from the approach of Einstein who, in a series of papers (1902, 1903, 1904) independently developed a formalism closely related to that of Gibbs, but used it as a probe to obtain empirical tests for the molecular/atomic hypothesis (cf. [Gearhart, 1990; Navarro, 1998; Uffink, 2006]).

and

What we know about a body can generally be described most accurately and most simply by saying that it is one taken at random from a great number (ensemble) of bodies which are completely described. (p. 163)

Note that Gibbs is somewhat non-committal about any particular interpretation of probability. (Of course, most of the presently distinguished interpretations of probability were only elaborated since the 1920s, and we cannot suppose Gibbs to have pre-knowledge of those distinctions.) A modern frequentist (for whom a probability of an event is the frequency with which that event occurs in a long sequence of similar cases) will have no difficulty with Gibbs' reference to an ensemble, and will presumably identify that notion with von Mises' notion of a *Kollektiv*. On the other hand, authors like Jaynes who favour a subjectivist interpretation of probability (in which the probability of an event is understood as a state of knowledge or belief about that event) have emphasized that in Gibbs' approach the ensemble is merely 'imagined' and a tool for representing our knowledge.

The ensemble is usually presented in the form of a probability density function ρ over Γ , such that $\int_A \rho(x) dx$ is the relative number of systems in the ensemble whose microstate $x = (\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N)$ lies in the region A . The evolution of an ensemble density ρ_0 at time $t = 0$ is dictated by the Hamiltonian equations of motion. In terms of the (formal) time evolution operator T_t , we get

$$(68) \quad \rho_t(x) = \rho_0(T_{-t}x)$$

or, in differential form:

$$(69) \quad \frac{\partial \rho_t(x)}{\partial t} = \{H, \rho\}$$

where $\{\cdot, \cdot\}$ denotes the Poisson bracket:

$$(70) \quad \{H, \rho\} = \sum_{i=1}^N \frac{\partial H}{\partial \vec{q}_i} \frac{\partial \rho}{\partial \vec{p}_i} - \frac{\partial H}{\partial \vec{p}_i} \frac{\partial \rho}{\partial \vec{q}_i}$$

A case of special interest is that in which the density function is stationary, i.e.

$$(71) \quad \forall t: \quad \frac{\partial \rho_t(x)}{\partial t} = 0.$$

This is what Gibbs calls the condition of *statistical equilibrium*. Gibbs notes that any density which can be written as a function of the Hamiltonian is stationary, and proceeds to distinguish special cases, of which the most important are:

$$(72) \quad \rho_E(x) = \frac{1}{\omega(E)} \delta(H(x) - E) \quad (\text{microcanonical})$$

$$(73) \quad \rho_\theta(x) = \frac{1}{Z(\theta)} \exp(-H(x)/\theta) \quad (\text{canonical})$$

$$(74) \quad \rho_{\theta, \alpha}(x, N) = \frac{1}{N! Z(\theta, \alpha)} \exp(-H(x)/\theta + \alpha N) \quad (\text{grand-canonical})$$

where $\omega(E)$, $Z(\theta)$ and $Z(\theta, \alpha)$ are normalization factors. In the following I will mainly discuss the canonical and microcanonical ensembles.

5.1 Thermodynamic analogies for statistical equilibrium

As indicated by the subtitle of the book, Gibbs' main goal was to provide a 'rational foundation' for thermodynamics. He approaches this issue quite cautiously, by pointing out certain analogies between relations holding for the canonical and microcanonical ensembles and results of thermodynamics. At no point does Gibbs claim to have reduced thermodynamics to statistical mechanics.

The very first analogy noticed by Gibbs is in the case of two systems, A and B put into thermal contact. This is modeled in statistical mechanics by taking the product phase space, $\Gamma_{AB} = \Gamma_A \times \Gamma_B$, and a Hamiltonian $H_{AB} = H_A + H_B + H_{\text{int}}$. If both A and B are described by canonical ensembles and if H_{int} is 'infinitely small' compared to the system Hamiltonian, then the combined system will be in statistical equilibrium if $\theta_A = \theta_B$. This, he says, "is entirely analogous to ... the corresponding case in thermodynamics" where "the most simple test of the equality of temperature of two bodies is that they remain in thermal equilibrium when brought into thermal contact" (ibid. p. 37). Clearly, Gibbs invites us to think of statistical equilibrium as analogous to thermal equilibrium, and θ as the analogue of the temperature of the system.⁴⁸

A second point of analogy is in reproducing the 'fundamental equation' (23) of thermodynamics:

$$(75) \quad dU = TdS + \sum_i F_i da_i$$

where a_i are the so-called external parameters (e.g. volume) and F_i the associated generalized forces (e.g. minus the pressure). For the canonical ensemble, Gibbs derives a relation formally similar to the above fundamental equation:⁴⁹

$$(76) \quad d\langle H \rangle = \theta d\sigma - \sum_i \langle A_i \rangle da_i.$$

Here, $\langle H \rangle$ is the expectation value of the Hamiltonian in the canonical ensemble, θ the modulus of the ensemble, σ the so-called Gibbs entropy of the canonical distribution:

$$(77) \quad \sigma[\rho_\theta] = - \int \rho_\theta(x) \ln \rho_\theta(x) dx,$$

a_i are parameters in the form of the Hamiltonian and the $\langle A_i \rangle = \langle \frac{\partial H}{\partial a_i} \rangle$ represent

⁴⁸A more elaborate discussion of the properties of the parameter θ and their analogies to temperature, is in Einstein (1902). That discussion also addresses the transitivity of thermal equilibrium, i.e. the Zeroth Law of thermodynamics (cf. paragraph 2).

⁴⁹See [Uhlenbeck and Ford, 1963; van Lith, 2001b] for details.

the ‘generalized forces’.⁵⁰ The equation suggests that the canonical ensemble averages might serve as analogues of the corresponding thermodynamic quantities, and θ and σ as analogues of respectively temperature and entropy.⁵¹

Note the peculiarly different role of θ and σ in (76): these are not expectations of phase space functions, but a parameter and a functional of the ensemble density ρ_θ . This has a significant conceptual implication. The former quantities may be thought of as averages, taken over the ensemble of some property possessed by each individual system in the ensemble. But for temperature θ and entropy σ , this is not so. In the case of θ one can diminish this contrast — at least when H is the sum of a kinetic and a potential energy term and the kinetic part is quadratic in the momenta, i.e. $H = \sum_i \alpha_i p_i^2 + U(q_1, \dots, q_n)$ — because of the well-known equipartition theorem. This theorem says that θ equals twice the expected kinetic energy for each degree of freedom:

$$(78) \quad \frac{\theta}{2} = \alpha_i \langle p_i^2 \rangle_\theta.$$

Thus, in this case, one can find phase functions whose canonical expectation values are equal to θ , and regard the value of such a function as corresponding to the temperature of an individual system.⁵² But *no* function χ on phase space exists such that

$$(79) \quad \sigma[\rho_\theta] = \langle \chi \rangle_\theta \quad \text{for all } \theta.$$

Thus, the Gibbs entropy cannot be interpreted as an average of some property of the individual members of the ensemble.

The next question is whether a differential equation similar to (76) can be obtained also for the microcanonical ensemble. In this case, it is natural to consider the same expressions $\langle A_i \rangle$ and $\langle H \rangle$ as above, but now taken as expectations with respect to the microcanonical ensemble, so that obviously $\langle H \rangle_{\text{mc}} = E$. The problem is then to find the microcanonical analogies to T and S . [Gibbs, 1902, p. 124–128, 169–171] proposes the following:

$$(80) \quad T \longleftrightarrow \left(\frac{d \ln \Omega(E)}{dE} \right)^{-1},$$

$$(81) \quad S \longleftrightarrow \ln \Omega(E),$$

where

$$(82) \quad \Omega(E) := \int_{H(x) \leq E} dp_1 \dots dq_n$$

⁵⁰A more delicate argument is needed if one wishes to verify that $-\langle \frac{\partial H}{\partial V} \rangle$ can really be identified with pressure, i.e. the average force per unit area on the walls of the container. Such an argument is given by [Martin-Löf, 1979, p. 21–25]

⁵¹A crucial assumption in this derivation is that the differential expressions represent infinitesimal elements of quasistatic processes during which the probability density always retains its canonical shape. This assumption is in conflict with a dynamical evolution [van Lith, 2001b, p. 141].

⁵²For proposals of more generally defined phase functions that can serve as an analogy of temperature, see [Rugh, 1997; Jepps *et al.*, 2000].

is known as the integrated structure function.

Remarkably, in a later passage, [Gibbs, 1902, p. 172–178] also provides a second pair of analogies to temperature and entropy, namely:

$$(83) \quad T \longleftrightarrow \left(\frac{d \ln \omega(E)}{dE} \right)^{-1}$$

$$(84) \quad S \longleftrightarrow \ln \omega(E),$$

where ω is the structure function

$$\omega(E) = \frac{d\Omega(E)}{dE} = \int_{H(x)=E} dx.$$

For this choice, the relation (75) is again reproduced. Thus, there appears to be a variety of choice for statistical mechanical quantities that may serve as thermodynamic analogue. Although Gibbs discussed various pro's and con's of the two sets, — depending on such issues as whether we regard the energy or the temperature as an independent variable, and whether we prefer expected values of most probable values — he does not reach a clear preference for one of them. (As he put it, system (80,81) is the more natural, while system (83,84) is the simpler of the two.) Still, Gibbs argued (*ibid.*, p. 183) that the two sets of analogies will approximately coincide for a very large number degrees of freedom. Nevertheless, this means there remains an underdetermination in his approach that one can hope to avoid only in the thermodynamic limit.

The expressions (81) and (84) are also known as the ‘volume entropy’ and the ‘surface entropy’. In modern textbooks the latter choice has been by far the most popular, perhaps because it coincides with the Gibbs entropy for the microcanonical ensemble: $\sigma[\rho_E] = \ln \omega(E)$. However, it has been pointed out that there are also general theoretical reasons to prefer the volume entropy (81), in particular because it is, unlike the surface entropy, an adiabatic invariant (see [Hertz, 1910; Rugh, 2001; Campisi, 2005]).

Of course, all of this is restricted to (statistical) equilibrium. In the case of non-equilibrium, one would obviously like to obtain further thermodynamical analogies that recover the approach to equilibrium (the ‘Minus First Law’, cf. p. 939) and an increase in entropy for adiabatic processes that start and end in equilibrium, or even to reproduce the kinetic equations on a full statistical mechanical basis. What Gibbs had to say on such issues will be the subject of the paragraphs 5.3 and 5.4.

But Gibbs also noted that a comparison of temperature and entropy with their analogies in statistical mechanics “would not be complete without a consideration of their differences with respect to units and zeros and the numbers used for their numerical specification” [Gibbs, 1902, p.183]. This will be taken up below in §5.2.

5.2 Units, zeros and the factor $N!$

The various expressions Gibbs proposed as analogies for entropy, i.e. (77,81,84), were presented without any discussion of ‘units and zeros’, i.e. of their physical dimension and the constants that may be added to these expressions. This was only natural because Gibbs singled out those expressions for their formal merit of reproducing the fundamental equation, in which only the combination TdS appears. He discussed the question of the physical dimension of entropy by noting that the fundamental equation remains invariant if we multiply the analogue for temperature — i.e. the parameter θ in the canonical case, or the functions (80 or (83) for the microcanonical case — by some constant K and the corresponding analogues for entropy — (77), (81) and (84) — by $1/K$. Applied to the simple case of the monatomic ideal gas of N molecules, he concluded that, in order to equate the analogues of temperature to the ideal gas temperature, $1/K$ should be set equal to

$$(85) \quad \frac{1}{K} = \frac{2}{3} \frac{c_V}{N},$$

where c_V is the specific heat at constant volume. He notes that “this value had been recognized by physicists as a constant independent of the kind of monatomic gas considered” [Gibbs, 1902, p. 185]. Indeed, in modern notation, $1/K = k$, i.e. Boltzmann’s constant.

Concerning the question of ‘zeros’, Gibbs noted that all the expressions proposed as analogy of entropy had the dimension of the logarithm of phase space volume and are thus affected by the choice of our units for length mass and time in the form of some additional constant (cf. [Gibbs, 1902, p. 19,183]). But even if some choice for such units is fixed, further constants could be added to the statistical analogs of entropy, i.e. arbitrary expressions that may depend on anything not varied in the fundamental equation. However, their values would disappear when differences of entropy are compared. And since only entropy differences have physical meaning, a question of determining these constants would thus appear to be immaterial. However, Gibbs went on to argue that “the principle that the entropy of any body has an arbitrary additive constant is subject to limitations when different quantities of the same substance are compared” [Gibbs, 1902, p. 206]. He formulated further conditions on how the additive constant may depend on the number N of particles in his final chapter.

Gibbs starts this investigation by raising the following problem. Consider the phase (i.e. microstate) $(\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N)$ of an N -particle system where the particles are said to be “indistinguishable”, “entirely similar” or “perfectly similar”.⁵³ Now, if we perform a permutation on the particles of such a system, should we regard the result as a different phase or not? Gibbs first argues that it “seems in accordance with the spirit of the statistical method” to regard such phases as the

⁵³Presumably, these terms mean (at least) that the Hamiltonian is invariant under their permutation, i.e. they have equal mass and interact in exactly the same way.

same. It might be urged, he says, that for such particles no identity is possible except that of qualities, and when comparing the permuted and unpermuted system, “nothing remains on which to base the identification of any particular particle of the first system with any particular particle of the second” [Gibbs, 1902, p. 187].

However, he immediately rejects this argument, stating that all this would be true for systems with “simultaneous objective existence”, but hardly applies to the “creations of the imagination”. On the contrary, Gibbs argues:

“The perfect similarity of several particles of a system will not in the least interfere with the identification of a particular particle in one case and with a particular particle in another. The question is one to be decided in accordance with the requirements of practical convenience in the discussion of the problems with which we are engaged” [Gibbs, 1902, p. 188]

He continues therefore by exploring both options, calling the viewpoint in which permuted phases are regarded as identical the *generic* phase, and that in which they are seen as distinct the *specific* phase. In modern terms the generic phase space is obtained as the quotient space of the specific phase space obtained by identifying all phase points that differ by a permutation (see [Leinaas and Myrheim, 1977]). In general, there are $N!$ different permutations on the phase of a system of N particles,⁵⁴ and there are thus $N!$ different specific phases corresponding to one generic phase. This reduces the generic phase space measure by an overall factor of $\frac{1}{N!}$ in comparison to the specific phase space. Since the analogies to entropy all have a dimension equal to the logarithm of phase space measure, this factor shows up as an further additive constant to the entropy, namely $-\ln N!$ in comparison to an entropy calculated from the specific phase. Gibbs concludes that when N is constant, “it is therefore immaterial whether we use [the generic entropy] or [the specific entropy], since this only affects the arbitrary constant of integration which is added to the entropy [Gibbs, 1902, p. 206].⁵⁵

However, Gibbs points out that this is *not* the case if we compare the entropies of systems with different number of particles. For example, consider two identical gases, each with the same energy U , volume V and number of particles N , in contiguous containers, and let the entropy of each gas be written as $S(U, V, N)$. Gibbs puts the entropy of the total system equal to the sum of the entropies:

$$(86) \quad S_{\text{tot}} = 2S(U, V, N).$$

Now suppose a valve is opened, making a connection between the two containers. Gibbs says that “we do not regard this as making any change in the entropy, although the gases diffuse into one another, and this process would increase the entropy if the gases were different” [Gibbs, 1902, p. 206-7]. Therefore, the entropy in this new situation is

⁵⁴This assumes that the molecular states \vec{p}_i, \vec{q}_i of the particles do not coincide. However the points in specific phase space for which one or more molecular states do coincide constitute a set of Lebesgue measure zero anyway.

⁵⁵The same conclusion also obtains for the Boltzmann entropy (61) [Huggett, 1999].

$$(87) \quad S'_{\text{tot}} = S_{\text{tot}}.$$

But the new system, is a gas with energy $2U$, volume $2V$, and particle number $2N$. Therefore, we obtain:

$$(88) \quad S'_{\text{tot}} = S(2U, 2V, 2N) = 2S(U, V, N),$$

where the right-hand side equation expresses the *extensivity* of entropy. This condition is satisfied (at least for large N) by the generic entropy but not by the specific entropy. Gibbs concludes “it is evident therefore that it is equilibrium with respect to generic phases, and not that with respect to specific, with which we have to do in the evaluation of entropy, . . . except in the thermodynamics of bodies in which the number of molecules of the various kinds is constant” [Gibbs, 1902, p. 207].

The issue expressed in these final pages is perhaps the most controversial in Gibbs’ book; at least it has generated much further discussion. Many later authors have argued that the insertion of a factor $1/N!$ in the phase space measure is obligatory to obtain “correct” results and, ultimately due to a lack of any metaphysical identity or “haecceity” of the perfectly similar particles considered. Some have even gone on to argue that quantum mechanics is needed to explain this. For example, [Huang, 1987, p. 154] writes “It is not possible to understand classically why we must divide [...] by $N!$ to obtain the correct counting of states. The reason is inherently quantum mechanical . . .”. However, many others deny this [Becker, 1967; van Kampen, 1984; Ray, 1984]. It would take me too far afield to discuss the various views and widespread confusion on this issue.

Let it suffice to note that Gibbs rejected arguments from the metaphysics of identity for the creations of the imagination. (I presume this may be taken to express that the phases of an N -particles system are theoretical constructs, rather than material objects.) Further, Gibbs did not claim that the generic view was correct and the specific view of incorrect; he preferred to settle the question by “practical convenience”. There are indeed several aspects of his argument that rely on assumptions that may be argued to be conventional. for example the ‘additivity’ demand (86) could be expanded to read more fully:

$$(89) \quad S_{\text{tot}}(U_1, V_1, N_1; U_2, V_2, N_2) + K_{\text{tot}} = S_1(U_1, V_1, N_1) + K_1 + S_2(U_2, V_2, N_2) + K_2,$$

Applied to the special case where S_1 and S_2 are identical functions taken at the same values of their arguments. The point to note here is that this relation only leads to (86) if we also employ the conventions $K_{\text{tot}} = K_1 + K_2$ and $K_1 = K_2$. Also, his cautious choice of words concerning (87) — “we do not regard this as making any change” — suggest that he wants to leave open whether this equation expresses a fact or a conventional choice on our part. But by and large, it seems fair to say that Gibbs’ criterion for practical convenience is simply the recovery of the properties usually assumed to hold for thermodynamic entropy.

As a final remark, note that the contrast mentioned here in passing by Gibbs, i.e. that in thermodynamics the mixing of identical gases, by allowing them to

diffuse into one another, does not change the entropy, whereas this process does increase entropy if the gases are different, implicitly refers to an earlier discussion of this issue in his 1875 paper [Gibbs, 1906, pp. 165–167]. The contrast between the entropy of mixing of identical fluids and that of different fluids noted on that occasion is now commonly known as the *Gibbs paradox*. (More precisely, this ‘paradox’ is that the entropy of mixing different fluids is a constant ($kT \ln 2$ in the above case) as long as the substances are different, and vanishes abruptly when they are perfectly similar; thus negating the intuitive expectation one might have had that the entropy of mixing should diminish gradually when the substances become more and more alike). Now note that in the the specific view, mixing different substances and mixing identical substances both lead to an entropy increase: in that view there is no Gibbs paradox, since there is no abrupt change when the substances become more and more alike. On the other hand, the adoption of the generic view, i.e. the division of the phase space measure by $N!$, is used by Gibbs to recover the usual properties of thermodynamic entropy *including* the Gibbs paradox — the discontinuity between mixing of different and identical gases.

Still, many authors seem to believe that the division by $N!$ is a procedure that *solves* the Gibbs paradox. But this is clearly not the case; instead, it is the specific viewpoint that avoids the paradox, while the generic viewpoint recovers the Gibbs paradox for the statistical mechanical analogies to entropy. The irony of it all is that, in statistical mechanics, the term “Gibbs paradox” is sometimes used to mean or imply the *absence* of the original Gibbs paradox in the specific point of view, so that a resolution of *this* “Gibbs paradox” requires the return of the original paradox.

5.3 Gibbs on the increase of entropy

As we have seen, the Gibbs entropy may be defined as a functional on arbitrary probability density functions ρ on phase space Γ :⁵⁶

$$(90) \quad \sigma[\rho] = - \int \rho(x) \ln \rho(x) dx$$

This expression has many well-known and useful properties. For example, under all probability densities restricted to the energy hypersurface $H(x) = E$, the microcanonical density (72) has the highest entropy. Similarly, one can show that of all distributions ρ with a given expectation value $\langle H \rangle_\rho$, the canonical distribution (73) has the highest entropy, and that of all distributions for which both $\langle H \rangle$ and $\langle N \rangle$ are given, the grand-canonical ensemble has the highest entropy.

⁵⁶Gibbs actually does not use the term entropy for this expression. He calls the function $\ln \rho$ the “index of probability”, and $-\sigma$ “the average index of probability”. As we have seen, Gibbs proposed more than one candidate for entropy in the microcanonical ensemble, and was well aware that: “[t]here may be [...], and there are, other expressions that may be thought to have some claim to be regarded as the [...] entropy with respect to systems of a finite number of degrees of freedom” [Gibbs, 1902, p. 169].

But suppose that ρ is not stationary. It will therefore evolve in the course of time, as given by $\rho_t(x) = \rho(T_{-t}x)$. One might ask whether this entropy will increase in the course of time. However, Liouville's theorem implies immediately

$$(91) \quad \sigma[\rho_t] = \sigma[\rho_0].$$

In spite of the superficial similarity to Boltzmann's H , the Gibbs entropy thus remains constant in time. The explanation of the Second Law, or an approach to equilibrium, cannot be so simple.

However, Gibbs warns us to proceed with great caution. Liouville's theorem can be interpreted as stating that the motion of ρ_t can be likened to motion in a (multidimensional) incompressible fluid. He thus compared the evolution of ρ to that of the stirring of a dye in an incompressible medium [Gibbs, 1902, p. 143-151]. In this case too, the average density of the dye, as well as the average of any function of its density, does not change. Still, it is a familiar fact of experience that by stirring tends to bring about a uniform mixture, or a state with uniform density, for which the expression $-\int \rho \ln \rho dx$ would have increased to attain its maximum value.

Gibbs saw the resolution of this contradiction in the definition of the notion of density. This, of course, is commonly taken as the limit of the quantity of dye in a spatial volume element, when the latter goes to zero. If we apply this definition, i.e. take this limit first, and then consider the stirring motion, we will arrive at the conclusion that $-\int \rho \ln \rho dx$ remains constant. But if we consider the density defined for a fixed finite (non-zero) volume element, and then stir for an indefinitely long time, the density may become 'sensibly' uniform, a result which is not affected if we subsequently let the volume elements become vanishingly small. The problem, as Gibbs saw it, is therefore one of the order in which we proceed to take two limits.

Gibbs was aware that not all motions in phase space produce this tendency toward statistical equilibrium, just as not every motion in an incompressible fluid stirs a dye to a sensibly homogeneous mixture. Nevertheless, as he concluded tentatively, "We might perhaps fairly infer from such considerations as have been adduced that an approach to a limiting condition of statistical equilibrium is the general rule, when the initial condition is not of that character" [Gibbs, 1902, p. 148].

5.4 Coarse graining

The most common modern elaboration of Gibbs' ideas is by taking recourse to a partitioning of phase space in cells, usually called "*coarse graining*". Instead of studying the original distribution function $\rho(x)$ we replace $\rho(x)dx$ by its phase average over each cell, by the mapping:

$$(92) \quad \mathcal{CG} : \rho(x) \mapsto \mathcal{CG}\rho(x) = \sum_i \hat{\rho}(i) \mathbf{1}_{\omega_i}(x),$$

where

$$(93) \quad \hat{\rho}(i) := \frac{\int_{\omega_i} \rho(x) dx}{\int_{\omega_i} dx},$$

and $\mathbf{1}$ denotes the characteristic function:

$$(94) \quad \mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{elsewhere.} \end{cases}$$

The usual idea is that such a partition matches the finite precision of our observational capabilities, so that a “coarse grained” distribution might be taken as a sufficient description of what is observable. Obviously, the average value of any function on Γ that does not vary too much within cells is approximately the same, whether we use the fine-grained or the coarse-grained distribution.

For any ρ one can also define the coarse grained entropy $\Sigma[\rho]$ as the composition of (92) and (90):

$$(95) \quad \Sigma[\rho] := \sigma[\mathcal{CG}\rho].$$

This coarse grained entropy need not be conserved in time. Indeed, it is easy to show (cf. [Tolman, 1938, p. 172]). that:

$$(96) \quad \Sigma[\rho] \geq \sigma[\rho].$$

Hence, if we assume that at some initial time that $\rho_0 = \mathcal{CG}\rho_0$, e.g. if $\rho_0 \propto \frac{1}{V_i} \mathbf{1}_{\omega_i}$ for some cell i , then for all t :

$$(97) \quad \Sigma[\rho_t] \geq \sigma[\rho_t] = \sigma[\rho_0] = \Sigma[\rho_0].$$

However, this does not imply that $\Sigma[\rho_t]$ is non-decreasing or that it approaches a limiting value as $t \rightarrow \infty$.

If a property, similar to the stirring of a dye holds for the dynamical evolution of ρ_t , one may have

$$(98) \quad \lim_{t \rightarrow \infty} \Sigma[\rho_t] = \Sigma[\rho_{mc}]$$

and hence, an approach towards equilibrium could emerge on the coarse-grained level. This convergence will of course demand a non-trivial assumption about the dynamics. In modern work this assumption is that the system has the *mixing* property (see paragraph 6.1).

5.5 Comments

Gibbs’ statistical mechanics has produced a formalism with clearly delineated concepts and methods, using only Hamiltonian mechanics and probability theory. It can and is routinely used to calculate equilibrium properties of gases and other systems by introducing a specific form of the Hamiltonian. The main problems that Gibbs has left open are, first, the motivation for the special choice of the

equilibrium ensembles and, second, that the quantities serving as thermodynamic analogies are not uniquely defined. However, much careful work has been devoted to show that, under certain assumptions about tempered interaction of molecules, unique thermodynamic state functions, with their desired properties are obtained in the ‘thermodynamic limit’ (cf. §6.3).

1. Motivating the choice of ensemble. While Gibbs had not much more to offer in recommendation of these three ensembles than their simplicity as candidates for representation for equilibrium, modern views often provide an additional story. First, the microcanonical ensemble is particularly singled out for describing an ensemble of systems in thermal isolation with a fixed energy E .

Arguments for this purpose come in different kinds. As argued by Boltzmann (1868), and shown more clearly by Einstein (1902), the microcanonical ensemble is the *unique* stationary density for an isolated ensemble of systems with fixed energy, if one assumes the ergodic hypothesis. Unfortunately, for this argument, the ergodic hypothesis is false for any system that has a phase space of dimension 2 or higher (cf. paragraph 6.1).

A related but more promising argument relies on the theorem that the measure P_{mc} associated with the microcanonical ensemble via $P_{\text{mc}}(A) = \int_A \rho_{\text{mc}}(x) dx$ is the unique stationary measure under all measures that are absolutely continuous with respect to P_{mc} , if one assumes that the system is metrically transitive (again, see paragraph 6.1).

This argument is applicable for more general systems, but its conclusion is weaker. In particular, one would now have to argue that physically interesting systems are indeed metrically transitive, and why measures that are not absolutely continuous with respect to the microcanonical one are somehow to be disregarded. The first problem is still an open question, even for the hard-spheres model (as we shall see in paragraph 6.1). The second question can be answered in a variety of ways.

For example, [Penrose, 1979, p. 1941] adopts a principle that every ensemble should be representable by a (piecewise) continuous density function, in order to rule out “physically unreasonable cases”. (This postulate implies absolute continuity of the ensemble measure with respect to the microcanonical measure by virtue of the Radon-Nikodym theorem.) See [Kurth, 1960, p. 78] for a similar postulate. Another argument, proposed by [Malament and Zabell, 1980], assumes that the measure P associated with a physically meaningful ensemble should have a property called ‘translation continuity’. Roughly, this notion means that the probability assigned to any measurable set should be a continuous function under small displacements of that set within the energy hypersurface. Malament & Zabell show that this property is equivalent to absolute continuity of P with respect to μ_{mc} , and thus singles out the microcanonical measure uniquely if the system is metrically transitive (see [van Lith, 2001b, for a more extensive discussion]).

A third approach, due to Tolman and Jaynes, more or less postulates the microcanonical density, as a appropriate description of our knowledge about the

microstate of a system with given energy (regardless of whether the system is metrically transitive or not).

Once the microcanonical ensemble is in place as a privileged description of an isolated system with a fixed energy, one can motivate the corresponding status for the other ensembles with relatively less effort. The canonical distribution is shown to provide the description of a small system S_1 in weak energetic contact with a larger system S_2 , acting as a ‘heat bath’ (see [Gibbs, 1902, p. 180–183]). Here, it is assumed that the total system is isolated and described by a microcanonical ensemble, where the total system has a Hamiltonian $H_{\text{tot}} = H_1 + H_2 + H_{\text{int}}$ with $H_2 \gg H_1 \gg H_{\text{int}}$. More elaborate versions of such an argument are given by Einstein (1902) and Martin-Löf (1979). Similarly, the grand-canonical ensemble can be derived for a small system that can exchange both energy and particles with a large system. (see [van Kampen, 1984]).

2. The ‘equivalence’ of ensembles. It is often argued in physics textbooks that the choice between these different ensembles (say the canonical and micro-canonical) is deprived of practical relevance by a claim that they are all “equivalent”. (See [Lorentz, 1916, p. 32] for perhaps the earliest version of this argument, or [Thompson, 1972, p. 72; Huang, 1987, p. 161–2] for recent statements.) What is meant by this claim is that if the number of constituents increases, $N \rightarrow \infty$, and the total Hamiltonian is proportional to N , the thermodynamic relations derived from each of them will coincide in this limit.

However, these arguments should not be mistaken as settling the *empirical* equivalence of the various ensembles, even in this limit. For example, it can be shown that the microcanonical ensemble admits the description of certain metastable thermodynamic states, (e.g. with negative heat capacity) that are excluded in the canonical ensemble (see [Touchette, 2003; Touchette et al., 2004, and literature cited therein]).

3. The coarse-grained entropy. The coarse-graining approach is reminiscent of Boltzmann’s construction of cells in his (1877b); cf. the discussion in paragraph 4.4). The main difference is that here one assumes a partition on phase-space Γ , where Boltzmann adopted it in the μ -space. Nevertheless, the same issues about the origin or status of a privileged partition can be debated (cf. p. 977). If one assumes that the partition is intended to represent what we *know* about the system, i.e. if one argues that all we know is whether its state falls in a particular cell ω_i , it can be argued that the its status is subjective. If one argues that the partition is meant to represent limitations in the precision of human observational possibilities, perhaps enriched by instruments, i.e. that we cannot observe more about the system than that its state is in some cell ω_i , one might argue that its choice is objective, in the sense that there are objective facts about what a given epistemic community can observe or not. Of course, one can then still maintain that the status of the coarse-graining would then be anthropocentric (see also the discussion in §7.5). However, note that Gibbs himself did not argue for a prefer-

ential size of the cells in phase space, but for taking the limit in which their size goes to zero in a different order.

4. Statistical equilibrium. Finally, a remark about Gibbs' notion of equilibrium. This is fundamentally different from Boltzmann's 1877 notion of equilibrium as the macrostate corresponding to the region occupying the largest volume in phase space (cf. section 4.4). For Gibbs, statistical equilibrium can only apply to an ensemble. And since any given system can be regarded as belonging to an infinity of different ensembles, it makes no sense to say whether an individual system is in statistical equilibrium or not. In contrast, in Boltzmann's case, equilibrium can be attributed to a single system (namely if the microstate of that system is an element of the set $\Gamma_{\text{eq}} \subset \Gamma$). But it is not guaranteed to remain there for all times.

Thus, one might say that in comparison with the orthodox thermodynamical notion of equilibrium (which is both stationary and a property of an individual system) Boltzmann (1877b) and Gibbs each made an opposite choice about which aspect to preserve and which aspect to sacrifice. See [Uffink, 1996b; Callender, 1999; Lavis, 2005] for further discussions.

6 MODERN APPROACHES TO STATISTICAL MECHANICS

This section will leave the more or less historical account followed in the previous sections behind, and present a selective overview of some influential modern approaches to statistical physics. In particular, we focus on ergodic theory (§ 6.1–6.2), the theory of the thermodynamic limit §6.3, the work of Lanford on the Boltzmann equation (§6.4), and the BBGKY approach in §6.5.

6.1 Ergodic theory

When the Ehrenfests critically reviewed Boltzmann's and Gibbs' approach to statistical physics in their renowned Encyclopedia article 1912, they identified three issues related to the ergodic hypothesis.

1. The ambiguity in Boltzmann's usage of "probability" of a phase space region (as either the relative volume of the region or the relative time spent in the region by the trajectory of the system).
2. The privileged status of the microcanonical probability distribution or other probability distributions that depend only on the Hamiltonian.
3. Boltzmann's argument that the microstate of a system, initially prepared in a region of phase space corresponding to a non-equilibrium macrostate, should tend to evolve in such a way that its trajectory will spend an overwhelmingly large majority of its time inside the region of phase space corresponding to the equilibrium macrostate Γ_{eq} .

In all these three problems, a more or less definite solution is obtained by adopting the ergodic hypothesis. Thus, the Ehrenfests suggested that Boltzmann's answer to the above problems *depended* on the ergodic hypothesis. As we have seen, this is correct only for Boltzmann's treatment of issue (2) in his 1868a. The doubtful status of the ergodic hypothesis, of course, highlighted the unresolved status of these problems in the Ehrenfests' point of view.

In later works the "ergodic problem" has become more exclusively associated with the first issue on the list above, i.e., the problem of showing the equality of phase and time averages. This problem can be formulated as follows. Consider a Hamiltonian system and some function f defined on its phase space Γ . The (infinite) time average of f , for a system with initial state x_0 may be defined as:

$$(99) \quad \overline{f(x_0)} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_t x_0) dt$$

where T_t is the evolution operator. On the other hand, for an ensemble of systems with density $\rho_t(x)$, the ensemble average of f is

$$(100) \quad \langle f \rangle_t = \int f(x) \rho_t(x) dx.$$

The ergodic problem is the question whether, or under which circumstances, the time average and ensemble average are equal, i.e.: $\overline{f(x_0)} \stackrel{?}{=} \langle f \rangle_t$. Note that there are immediate differences between these averages. \overline{f} depends on the initial state x_0 , in contrast to $\langle f \rangle$. Indeed, each choice of an initial phase point gives rise to another trajectory in phase space, and thus gives, in general, another time average. Secondly, $\langle f \rangle$ will in general depend on time, whereas \overline{f} is time-independent. Hence, a general affirmative answer to the problem cannot be expected.

However, in the case of a stationary ensemble (statistical equilibrium) the last disanalogy disappears. Choosing an even more special case, the microcanonical ensemble ρ_{mc} , the simplest version of the ergodic problem is the question:

$$(101) \quad \overline{f(x_0)} \stackrel{?}{=} \langle f \rangle_{mc}.$$

Now it is obvious that if Boltzmann's ergodic hypothesis is true, i.e. if the trajectory of the system traverses all points on the energy hypersurface Γ_E , the desired equality holds. Indeed, take two arbitrary points x and y in Γ_E . The ergodic hypothesis implies that there is a time τ such that $y = T_\tau x$. Hence:

$$\begin{aligned} \overline{f(y)} &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_{t+\tau} x) dt \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \left(\int_0^\tau f(T_t x) dt + \int_0^T f(T_t x) dt \right) \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_t x) dt = \overline{f(x)} \end{aligned}$$

In other words, \bar{f} must be constant over Γ_E , and hence, also equal to the micro-canonical expectation value.

For later reference we note another corollary: the ergodic hypothesis implies that ρ_{mc} is the only stationary density on Γ_E (cf. section 4.1).

The Ehrenfests doubted the validity of the ergodic hypothesis, as Boltzmann had himself, and therefore proposed an alternative, which they called the *quasi-ergodic hypothesis*. This states that the trajectory lies dense in Γ_E , i.e., x_t will pass through every open subset in Γ_E , and thus come arbitrarily close to every point in Γ_E . The system may be called quasi-ergodic if this holds for all its trajectories. As we have seen, this formulation seems actually closer to what Boltzmann may have intended, at least in 1871, than his own literal formulation of the hypothesis.

Not long after the Ehrenfests' review, the mathematical proof was delivered that the ergodic hypothesis cannot hold if Γ_E is a more than one-dimensional manifold [Rosenthal, 1913; Plancherel, 1913]. The quasi-ergodic hypothesis, on the other hand, cannot be immediately dismissed. In fact, it may very well be satisfied for Hamiltonian systems of interest to statistical mechanics. Unfortunately, it has remained unclear how it may contribute to a solution to the ergodic problem. One might hope, at first sight, that for a quasi-ergodic system time averages and micro-canonical averages coincide for continuous functions, and that the micro-canonical density ρ_{mc} is the only continuous stationary density. But even this is unknown. It is known that quasi-ergodic systems may fail to have a unique stationary measure [Nemytskii and Stepanov, 1960, p. 392]. This is not to say that quasi-ergodicity has remained a completely infertile notion. In topological ergodic theory, the condition is known under the name of "minimality", and implies several interesting theorems (see [Petersen, 1983, p. 152ff]).

While the Rosenthal-Plancherel result seemed to toll an early death knell over ergodic theory in 1913, a unexpected revival occurred in the early 1930s. These new results were made possible by the stormy developments in mathematics and triggered by Koopman's results, showing how Hamiltonian dynamics might be embedded in a Hilbert space formalism where the evolution operators T_t are represented as a unitary group. This made a whole array of mathematical techniques (e.g. spectral analysis) available for a new attack on the problem.

The first result was obtained by von Neumann in a paper under the promising (but misleading) title "Proof of the Quasi-Ergodic Hypothesis" 1932. His theorem was strengthened by G.D. Birkhoff in a paper entitled "Proof of the Ergodic Theorem" 1931, and published even before von Neumann's.

Since their work, and all later work in ergodic theory, involves more precise mathematical notions, it may be worthwhile first to introduce a more abstract setting of the problem. An abstract *dynamical system* is defined as a tuple $\langle \Gamma, \mathcal{A}, \mu, T \rangle$, where Γ as an arbitrary set, \mathcal{A} is a σ -algebra of subsets of Γ , called the 'measurable' sets in Γ , and μ is a probability measure on Γ , and T denotes a one-parameter group of one-to-one transformations T_t on Γ (with $t \in \mathbb{R}$ or $t \in \mathbb{Z}$) that represent the evolution operators. The transformations T_t are assumed to be measure-preserving, i.e. $\mu(T_t A) = \mu(A)$ for all $A \in \mathcal{A}$. In the more concrete

setting of statistical mechanics, one may take Γ to be the energy hypersurface, \mathcal{A} the collection of its Borel subsets, μ the microcanonical probability measure and T the evolution induced by the Hamiltonian equations.

The von Neumann-Birkhoff ergodic theorem can be formulated as follows:

ERGODIC THEOREM: Let $\langle \Gamma, \mathcal{A}, \mu, T \rangle$ be any dynamical system and f be an integrable function on Γ . Then

- (i) $\overline{f(x)} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_t x) dt$ exists for almost all x ;
i.e. the set of states $x \in \Gamma$ for which $\overline{f(x)}$ does not exist has μ -measure zero.
- (ii) $\overline{f(x)} = \langle f \rangle_\mu$ for almost all x iff the system is metrically transitive.

Here, metric transitivity means that it is impossible to carve up Γ in two regions of positive measure such that any trajectory starting in one region never crosses into the other. More precisely:

METRIC TRANSITIVITY: A dynamical system is called metrically transitive⁵⁷ iff the following holds: for any partition of Γ into disjoint sets A_1, A_2 such that $T_t A_1 = A_1$ and $T_t A_2 = A_2$, it holds that $\mu(A_1) = 0$ or $\mu(A_2) = 0$.

It is not difficult to see why this theorem may be thought of as a successful solution of the original ergodic problem under a slight reinterpretation. First, metric transitivity captures in a measure-theoretic sense the idea that trajectories wander wildly across the energy hypersurface, allowing only exceptions for a measure zero set. Secondly, the theorem ensures the equality of time and microcanonical ensemble average, although only for integrable functions and, again, with the exception of a measure zero set. But that seemed good enough for the taste of most physicists.

The ergodic theorem was therefore celebrated as a major victory. In the words of Reichenbach:

Boltzmann introduced [...] under the name of *ergodic hypothesis* [...] the hypothesis that the phase point passes through every point of the energy hypersurface. This formulation is easily shown to be untenable. It was replaced by P. and T. Ehrenfest by the formulation that the path comes close to every point within any small distance ϵ which we select and which is greater than 0.

There still remained the question whether the ergodic hypothesis must be regarded as an independent supposition or whether it is derivable from the canonical equations, as Liouville's theorem is.

⁵⁷This name is somewhat unfortunate, since the condition has nothing to do with metric in the sense of distance, but is purely measure-theoretical. Metrically transitive systems are also called 'metrically indecomposable' or, especially in the later literature 'ergodic'. I will stick to the older name in order to avoid confusion with the ergodic hypothesis.

This problem[. . .] was finally solved through ingenious investigations by John von Neumann and George Birkhoff, who were able to show that the second alternative is true. [...] With von Neumann and Birkhoff's theorem, deterministic physics has reached its highest degree of perfection: the strict determinism of elementary processes is shown to lead to statistical laws for macroscopic occurrences." [Reichenbach, 1956, p. 78]

Unfortunately, nearly everything stated in this quotation is untrue.

Problems

1. Do metrically transitive systems exist? An immediate question is of course whether metrically transitive systems exist. In a mathematical sense of 'exist' the answer is affirmative. More interesting is the question of whether one can show metric transitivity for any model that is realistic enough to be relevant to statistical mechanics.

A few mechanical systems have been explicitly proven to be metrically transitive. For example: one hard sphere moving in a vessel with a convex scatterer, or a disc confined to move in a 'stadium' (two parallel line-segments connected by two half circles) or its three-dimensional analogue: one hard sphere moving in a cylinder, closed on both sides by half-spheres. But in statistical mechanics one is interested in systems with many particles.

In [1963], Sinai announced he had found a proof that a gas consisting of N hard spheres is metrically transitive. The ergodic theorem thus finally seemed to be relevant to physically interesting gas models. Of course, the hard-spheres-model is an idealization too, but the general expectation among physicists was that a transition to more sophisticated models of a gas system would only make the metric transitivity even more likely and plausible, even though admittedly harder to prove.

The problem proves to be extraordinarily tedious, and Sinai's proof was complicated and, actually, never completely published. But many partial results were. In fact, the development of ideas and techniques needed for the effort contributed much to the emergence of a vigorous mathematical theory, nowadays called 'ergodic theory'. And since Sinai's claim seemed so desirable, many books and articles presented the claim as a solid proven fact (e.g. [Lebowitz and Penrose, 1973; Sklar, 1993]).

But by the 1980s, the delay in the publication of a complete proof started to foster some doubts about the validity of the claim. Finally, [Sinai and Chernov, 1987, p. 185] wrote: "The announcement made in [[Sinai, 1963]] for the general case must be regarded as immature." What has been shown rigorously is that a system of three hard spheres is metrically transitive. Recently, the problem has been taken further by [Szász, 1996] and [Simányi and Szász, 1999]. They have ascertained that for a model of N hard spheres, the ergodic component, i.e. a

subset of the energy hypersurface on which the assumption of metric transitivity holds has positive measure. The full problem, however, still awaits solution.

2. Infinite times. In the definition of the time average (99) the limit $T \rightarrow \infty$ is taken. This brings along a number of problems:

- (i). The time average is interesting because it is experimentally accessible. The hope is that it represents the equilibrium value of f . But the limit $T \rightarrow \infty$ tells us nothing about what happens in a finite time. What is empirically accessible, at best, is the quantity $\frac{1}{T} \int_0^T f(T_t x_0) dt$ for a large but finite T . This expression can still deviate arbitrarily far from the limiting value.
- (ii). The limit may even exist while the system is not in equilibrium. A time-averaged value need not be an equilibrium value, because in general

$$(102) \quad \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T f(T_t x) dt \neq \lim_{t \rightarrow \infty} f(T_t x).$$

For periodical motions, for example, the left-hand side exists but the right-hand side does not.

- (iii) Empirically, equilibrium often sets in quite rapidly. But the time T needed to make $\frac{1}{T} \int_0^T f(T_t x_0) dt$ even remotely close to $\langle f \rangle_{\text{mc}}$ might be enormous, namely of the order of Boltzmann's estimate of the Poincaré-recurrence times! (See also [Jaynes, 1967, p. 94].)

3. The measure-zero problem. The result that the ergodic theorem provides is that for metrically transitive systems $\overline{f(x)} = \langle f \rangle_{\text{mc}}$ except for a set of microstates with measure zero. So the suggestion here is that this set of exceptions is in some sense negligible. And, as judged from the probability measure μ_{mc} , that is obviously true. But a set of measure zero need not be negligible in any other sense. It is well-known that if one compares 'smallness in measure' with other natural criteria by which one can judge the 'size' of sets, e.g. by their cardinality, dimension or Baire category, the comparisons do not match. Sets of measure zero can be surprisingly large by many other standards [Sklar, 1993, pp. 181–188].

More importantly, one might choose another measure μ' , such that μ -measure zero sets are no longer sets of μ' -measure zero and conversely. It is of course the choice of the measure that determines which sets have measure zero. Thus, if one decides to disregard or neglect sets with a microcanonical measure zero, a privileged status of the microcanonical measure is already presupposed. But this means the virtue of the ergodic theorem as a means of motivating a privileged role of the microcanonical measure is diminished to a self-congratulating one.

6.2 The mixing property, K systems and Bernoulli systems

Ergodic theory, the mathematical field that emerged from the theorems of Birkhoff and von Neumann, may be characterized as a rigorous exploration of the question to what extent a deterministic, time-reversal invariant dynamical system may give rise to random-like behaviour on a macroscopic scale, by assuming various special properties on its dynamics.

In its modern incarnation, this theory distinguishes a hierarchy of such properties that consists of various strengthenings of metric transitivity. Perhaps the most important are the mixing property, the property of being a ‘ K system’ and the Bernoulli systems. The higher up one goes this ladder, the more ‘random’ behaviour is displayed. The evolution at the microlevel is in all cases provided by the deterministic evolution laws. In the (extensive) literature on the subject, many more steps in the hierarchy are distinguished (such as ‘weak mixing’, ‘weak Bernoulli’, ‘very weak Bernoulli’ etc.), and also some properties that do not fit into a strict linear hierarchy (like the ‘Anosov’ property, which relies on topological notions rather than on a purely measure-theoretical characterization of dynamical systems). It falls beyond the scope of this paper to discuss them.

Mixing

The idea of *mixing* is usually attributed to Gibbs, in his comparison of the evolution of ensembles with stirring of a dye into an incompressible fluid (cf. section 5.4). Even if initially the fluid and the dye particles occupy separate volumes, stirring will eventually distribute the dye particles homogeneously over the fluid. The formal definition is:

MIXING: A dynamical system $\langle \Gamma, \mathcal{A}, \mu, T \rangle$ is called mixing iff $\forall A, B \in \mathcal{A}$

$$(103) \quad \lim_{t \rightarrow \infty} \mu(T_t A \cap B) = \mu(A)\mu(B).$$

In an intuitive sense the mixing property expresses the idea that the dynamical evolution will thoroughly stir the phase points in such a way that points initially contained in A eventually become homogeneously distributed over all measurable subsets B of Γ . One can easily show that *mixing* is indeed a stronger property than metric transitivity, by applying the condition to an invariant set A and choosing $B = A$. The converse statement does not hold. (E.g.: the one-dimensional harmonic oscillator is metrically transitive but not mixing).

Again, there is an interesting corollary in terms of probability measures or densities. Consider a mixing system, and a time-dependent probability density ρ_t , such that ρ_t is *absolutely continuous* with respect to the microcanonical measure μ . (This means that all sets $A \in \mathcal{A}$ with $\mu(A) = 0$, also have $\int_A \rho_t(x) dx = 0$, or equivalently, that ρ_t is a proper density function that is integrable with respect to μ .) In this case, the probability measure associated with ρ_t converges, as $t \rightarrow \infty$, to the microcanonical measure. Thus, an ensemble of mixing systems

with an absolutely continuous density will asymptotically approach to statistical equilibrium. Note that the same result will also hold for $t \rightarrow -\infty$, so that there is no conflict with the time reversal invariance. Is it in conflict with Poincaré's recurrence theorem? No, the recurrence theorem is concerned with microstates (phase points), and not probability densities. Even when almost all trajectories eventually return close by their original starting point, the recurrence time will differ for each phase point, so that the evolution of an ensemble of such points can show a definite approach to statistical equilibrium.

Note also that if the result were used as an argument for the privileged status of the microcanonical measure (viz., as the unique measure that all absolutely continuous probability distributions evolve towards), the strategy would again be marred by the point that the condition of absolute continuity already refers to the microcanonical measure as a privileged choice.

Despite the elegance of the mixing property, we can more or less repeat the critical remarks made in the context of the ergodic theorem. In the first place, the condition considers the limit $t \rightarrow \infty$, which implies nothing about the rate at which convergence takes place. Secondly, the condition imposed is trivially true if we choose A or B to be sets of measure zero. Thus, the mixing property says nothing about the behaviour of such sets during time evolution. And thirdly, one is still faced with the question whether the mixing property holds for systems that are physically relevant for statistical mechanics. And since the property is strictly stronger than metric transitivity, this problem is at least as hard.

K systems

The next important concept is that of a *K system* ('K' after Kolmogorov). For simplicity, we assume that time is discrete, such that $T_t = T^t$, for $t \in \mathbb{Z}$. There is a perfectly analogously defined concept for continuous time, called *K flows* (cf. Emch, this volume, Definition 10.3.2).

*K SYSTEM:*⁵⁸ A dynamical system $(\Gamma, \mathcal{A}, \mu, T)$ is called a *K system* if there is a subalgebra $\mathcal{A}_0 \subset \mathcal{A}$, such that

1. $T^n \mathcal{A}_0 \subset T^m \mathcal{A}_0$ for times $m < n$; where \subset denotes proper inclusion.
2. the smallest σ -algebra containing $\cup_{n=1}^{\infty} T^{-n} \mathcal{A}_0$ is \mathcal{A} .
3. $\cap_{n=1}^{\infty} T^n \mathcal{A}_0 = \mathcal{N}$, where \mathcal{N} is the σ -algebra containing only sets of μ -measure zero or one.

At first sight, this definition may appear forbiddingly abstract. One may gain some intuition by means of the following example. Consider a finite partition

⁵⁸There is a considerable variation in the formulation of this definition [Cornfeld *et al.*, 1982; Batterman, 1991; Berkovitz *et al.*, 2006]. The present formulation adds one more. It is identical to more common definitions if one replaces n and m in the exponents of T by $-n$ and $-m$ respectively.

$\alpha = \{A_1, \dots, A_m\}$ of Γ into disjoint cells and the so-called *coarse-grained history* of the state of the system with respect to that partition. That is, instead of the detailed trajectory x_t , we only keep a record of the labels i of the cell A_i in which the state is located at each instant of time, until time $t=0$:

$$(104) \quad \dots i_{-k}, \dots, i_{-3}, i_{-2}, i_{-1}, i_0 \quad i_{-k} \in \{1, \dots, m\}, \quad k \in \mathbb{N}.$$

This sequence is completely determined by the microstate x at $t = 0$:

$$(105) \quad i_{-k}(x) = \sum_{j=1}^m j \mathbf{1}_{A_j}(T^{-k}x)$$

where $\mathbf{1}$ denotes the characteristic function (94). Yet, as we shall see, for a K system, this sequence typically behaves in certain respects like a random sequence. Observe that

$$(106) \quad i_{-k}(x) = j \iff T^{-k}x \in A_j \iff x \in T^k A_j;$$

so we can alternatively express the coarse-grained history by means of evolutions applied to the cells in the partition. If $T\alpha := \{TA_1, \dots, TA_m\}$, let $\alpha \vee T\alpha := \{A_i \cup TA_j : i, j = 1, \dots, m\}$ denote the common refinement of α and $T\alpha$. Saying that x belongs to $A_i \cup TA_j$ is, of course, equivalent to providing the last two terms of the sequence (104). Continuing in this fashion, one can build the refinement

$$(107) \quad \bigvee_{k=0}^{\infty} T^k \alpha = \alpha \vee T\alpha \vee T^2\alpha \dots \vee T^k \alpha \vee \dots,$$

each element of which corresponds to a particular coarse-grained history (104) up to $t=0$. The collection (107) is no longer finite, but still a countable partition of Γ .

Now take \mathcal{A}_0 to be the σ -algebra generated from the partition $\bigvee_{k=0}^{\infty} T^k \alpha$. Clearly, the events in this algebra are just those whose occurrence is completely decided whenever the coarse-grained history is known. In other words, for all $A \in \mathcal{A}_0$, $\mu(A|C)$ is zero or one, if C is a member of (107). It is easy to see that $T^{-m}\mathcal{A}_0$ is just the σ -algebra generated from $T^{-m}\bigvee_{k=0}^{\infty} T^k \alpha = \bigvee_{k=-m}^{\infty} T^k \alpha$, i.e. from the partition characterizing the coarse-grained histories up to $t = m$. Since the latter partition contains the history up to $t = n$ for all $n < m$, we have:

$$(108) \quad T^{-m}\mathcal{A}_0 \subseteq T^{-n}\mathcal{A}_0 \text{ for all } n < m.$$

This is equivalent to condition 1, but with ‘ \subset ’ replaced by ‘ \subseteq ’.

Further, to explain condition 2, note that the smallest σ -algebra containing $\bigcup_{n=1}^N T^{-n}\mathcal{A}_0$ is generated by the union of the partitions $\bigvee_{k=-n}^{\infty} T^k \alpha$ for all $n \leq N$, which in view of (108) is just $T^{-N}\mathcal{A}_0$. Thus, condition 2 just says that if we extend the record of the coarse-grained history to later times $t = N > 0$, and let $N \rightarrow \infty$, the partition eventually becomes sufficiently fine to generate all measurable sets

in \mathcal{A} . This is a strong property of the dynamics. It means that the entire coarse-grained record, extending from $-\infty$ to ∞ , provides all information needed to separate all the measurable sets in \mathcal{A} , (except, possibly, if they differ by a measure zero set.)

Similarly, in order to explain condition 3, note that (108) implies that $\bigcap_{n=1}^N T^n \mathcal{A}_0 = T^N \mathcal{A}_0$, which is generated from $\bigvee_{k=0}^{\infty} T^k \alpha$, i.e., the coarse-grained histories up to time $-N$. Thus, condition 3 expresses the demand that, as we let $N \rightarrow \infty$, the class of events that are settled by the coarse-grained histories up to time $t = -N$ shrinks to the ‘trivial’ algebra of those sets that have probability one or zero. In other words, for every event $A \in \mathcal{A}$, with $0 < \mu(A) < 1$, the occurrence of A is undecided at some early stage of the coarse-grained history.

Yet the truly remarkable feature of K systems lies in the strict inclusion demanded in condition 1: at any time n , the collection of events decided by the coarse-grained histories up to n , is strictly smaller than the collection of events decided at time $n+1$. Since the latter is generated from the former by adding the partition $T^{-(n+1)}\alpha$ to the partition $\bigvee T^{-k}\alpha$, this means that at each time n the question which cell of the partition is occupied at time $n+1$ is not answerable from the knowledge of the previous coarse-grained history. This is quite a remarkable property for a sequence generated by a deterministic law of motion, although, of course, it is familiar for random sequences such as tosses with a die or spins of a roulette wheel.

In this attempt at elucidation, we have presupposed a particular finite partition α . One may ask whether there always is, for each Kolmogorov system, such a partition. The answer is yes, provided the system obeys some mild condition (that $\langle \Gamma, \mathcal{A}, \mu \rangle$ is a *Lebesgue space*).⁵⁹ Another question is whether the claims made about coarse-grained histories are specific for this particular partition. The answer is no. One may show that, given that they hold for some partition α , they also hold for *any* choice of a finite partition of Γ . (Very roughly speaking: because the partition $\bigvee_n T^n \alpha$ generates the σ -algebra of all events, the coarse-grained histories constructed from another finite partition can be reconstructed in terms of the coarse-grained histories in terms of α .)

Bernoulli systems

The strongest property distinguished in the ergodic hierarchy is that of *Bernoulli systems*. To introduce the definition of this type of dynamical systems, it is useful to consider first what is usually known as a ‘Bernoulli’ scheme. Consider an elementary chance set-up with outcomes $\{A_1, \dots, A_m\}$ and probabilities p_j . A Bernoulli scheme is defined as the probability space obtained from doubly infinite sequences of independent identically distributed repetitions of trials on this elementary set-up. Formally, a Bernoulli scheme for a set (or ‘alphabet’) $\alpha = \{1, \dots, m\}$ with

⁵⁹Roughly, this condition means that $\langle \Gamma, \mathcal{A}, \mu \rangle$ is isomorphic (in a measure-theoretic sense) to the interval $[0, 1]$, equipped with the Lebesgue measure. (See [Cornfeld *et al.*, 1982, p. 449] for the precise definition).

probabilities $\{p_j\}$ is the probability space $\langle \Gamma, \mathcal{A}, \mu \rangle$, where Γ is the set of all doubly infinite sequences

$$(109) \quad \eta = (\dots, i_{-2}, i_{-1}, i_0, i_1, i_2, \dots) \quad i_k \in \{1, \dots, m\}; k \in \mathbb{Z}$$

and \mathcal{A} is defined as the smallest σ -algebra on Γ containing the sets:

$$(110) \quad A_k^j := \{\eta \in \Gamma : i_k = j\}.$$

\mathcal{A} is also known as the cylinder algebra. Further, we require of a Bernoulli scheme that:

$$(111) \quad \mu(A_k^j) = p_j \text{ for all } k \in \mathbb{Z}.$$

One can turn this probability space into a dynamical system by introducing the discrete group of transformations T^m , $m \in \mathbb{Z}$, where T denotes the shift, i.e. the transformation on Γ that shifts each element of a sequence η one place to the left:

$$(112) \quad \text{For all } k \in \mathbb{Z}: \quad T(i_k) = i_{k-1}.$$

Thus we define:

BERNOULLI SYSTEM: A dynamical system $\langle \Gamma, \mathcal{A}, \mu, T \rangle$ with a discrete time evolution T is a Bernoulli-system iff there is a finite partition $\alpha = \{A_1, \dots, A_m\}$ of Γ such that the doubly infinite coarse-grained histories are (isomorphic to) a Bernoulli scheme for α with distribution

$$(113) \quad p_i = \mu(A_i) \quad i \in \{1, \dots, m\}.$$

Thus, for a Bernoulli system, the coarse-grained histories on α behave as randomly as independent drawings from an urn. These histories show no correlation at all, and the best prediction one can make about the location of the state at time $n + 1$, even if we know the entire coarse-grained history from minus infinity to time n , is no better than if we did not know anything at all. One can show that every Bernoulli-system is also a K-system, but that the converse need not hold.

Discussion

Ergodic theory has developed into a full-fledged mathematical discipline with numerous interesting results and many open problems (for the current state of the field, see [Cornfeld *et al.*, 1982; Petersen, 1983; Mañé, 1987]). Yet the relevance of the enterprise for the foundations of statistical mechanics is often doubted. Thus [Earman and Rédei, 1996] argue that the enterprise is not relevant for explaining ‘why phase averaging works’ in equilibrium statistical mechanics; [Albert, 2000, p. 70] even calls the effort poured into rigorous proofs of ergodicity “nothing more nor less — from the standpoint of foundations of statistical mechanics — than a waste of time”. (For further discussions, see: [Farquhar, 1964; Sklar, 1973; Friedman, 1976; Malament and Zabell, 1980; Leeds, 1989; van Lith, 2001a; Frigg, 2004; Berkovitz *et al.*, 2006])

This judgment is usually based on the problems already indicated above; i.e. the difficulties of ascertaining that even the lowest property on the ergodic hierarchy actually obtains for interesting physical models in statistical mechanics, the empirical inaccessibility of infinite time averages, and the measure zero problem. Also, one often appeals to the Kolmogorov-Arnold-Moser (KAM) results⁶⁰ in order to temper the expectations that ergodicity could be a generic property of Hamiltonian systems. These difficulties are serious, but they do not, in my opinion, justify a definitive dismissal of ergodic theory.

Instead, it has been pointed out by [Khinchin, 1949; Malament and Zabell, 1980; Pitowsky, 2001] that further progress may be made by developing the theory in conditions in which (i) the equality of ensemble averages and time averages need not hold for *all* integrable functions, but for only a physically motivated subclass, (ii) imposing conditions that fix the rate of convergence in the infinite time limits in (99) and (103) and (iii) relaxing the conditions on what counts as an equilibrium state. Indeed important progress concerning (i) has been achieved in the ‘theory of the thermodynamic limit’, described in paragraph 6.3. It is clear that further alterations may be mathematically obstreperous; and that any results that might be obtained will not be as simple and general as those of the existing ergodic theory. But there is no reason why progress in these directions should be impossible. See e.g. [Vranas, 1998; van Lith, 2001b].

The measure zero problem, I would argue, is unsolvable within any “merely” measure-theoretic setting of the kind we have discussed above. The point is, that any measure theoretic discussion of dynamical systems that differ only on measure zero sets are, in measure-theoretical terms, isomorphic and usually identified. Measure theory has no way of distinguishing measure zero sets from the empty set. Any attempt to answer the measure zero problem should call upon other mathematical concepts. One can expect further light only by endowing the phase space with further physically relevant structure, e.g. a topology or a symplectic form (cf. [Butterfield, 2006; Belot, 2006]).

Furthermore, even if ergodic theory has little of relevance to offer to the explanation of ‘why phase averaging works’ in the case of equilibrium statistical mechanics, this does not mean it is a waste of time. Recall that the equality of phase and time averages was only one of several points on which the Ehrenfests argued that claims by Boltzmann could be substantiated by an appeal to the ergodic hypothesis. Another point was his (1877) claim that a system initially in a non-equilibrium macrostate should tend to evolve towards the equilibrium macrostate.

⁶⁰Quite roughly, the KAM theorems show that some Hamiltonian systems for which trajectories are confined to an invariant set in phase space of small positive measure — and therefore *not* metrically transitive —, will continue to have that property when a sufficiently small perturbation is added to their Hamiltonian (for a more informative introduction, see [Tabor, 1989]). This conclusion spoils the (once common) hope that non-metrically transitive systems were rare and idealized exceptions among Hamiltonian systems, and that they could always be turned into a metrically transitive system by acknowledging a tiny perturbation from their environment. As we have seen (p. 958), Boltzmann (1868) had already expressed this hope for the ergodic hypothesis.

It is ironic that some critics of ergodic theory dismiss the attempt to show in what sense and under which conditions the microstate does display a tendency to wander around the entire energy hypersurface as irrelevant, while relying on a rather verbal and pious hope that this will “typically” happen without any dynamical assumption to fall back on. Clearly, the ergodic hierarchy might still prove relevant here.

Still, it is undeniable that many concrete examples can be provided of systems that are not ergodic in any sense of the word and for which equilibrium statistical mechanics should still work. In a solid, say an ice cube, the molecules are tightly locked to their lattice site, and the phase point can access only a minute region of the energy hypersurface. Similarly, for a vapour/liquid mixture in a \cap -shaped vessel in a gravity field, molecules may spend an enormously long proportion of time confined to the liquid at the bottom of one leg of the vessel, even though the region corresponding to being located in the other leg is dynamically accessible. And still one would like to apply statistical mechanics to explain their thermal properties.

Summing up, even admitting that ergodic theory cannot provide the whole story in all desired cases does not mean it is irrelevant. I would argue that, on a qualitative and conceptual level, one of the most important achievements of ergodic theory is that it has made clear that strict determinism on the microscopic level is not incompatible with random behaviour on a macroscopic level, even in the strong sense of a Bernoulli system. This implies that the use of models with a stochastic evolution like urn drawings, that Boltzmann used in 1877, or the dog flea model of the Ehrenfests, (cf. §7.2), are not necessarily at odds with an underlying deterministic dynamics.

6.3 Khinchin’s approach and the thermodynamic limit

In the ‘hard core’ version of ergodic theory, described in the previous two paragraphs, one focuses on abstract dynamical systems, i.e. the only assumptions used are about a measure space equipped with a dynamical evolution. It is not necessary that this dynamics arises from a Hamiltonian. Further, it is irrelevant in this approach whether the system has a large number of degrees of freedom. Indeed, the ‘baker transformation’, an example beloved by ergodic theorists because it provides a dynamical system that possesses *all* the properties distinguished in the ergodic hierarchy, uses the unit square as phase space, and thus has only two degrees of freedom. On the other hand, Hamiltonian systems with large numbers of degrees of freedom, may fail to pass even the lowest step of the ergodic hierarchy, i.e. metric transitivity.

This aspect of ergodic theory is often criticized, because the thermal behaviour of macroscopic systems that the foundations of statistical mechanics ought to explain, arguably appears only when their number of degrees of freedom is huge. As Khinchin puts it:

All the results obtained by Birkhoff and his followers [...] pertain to

the most general type of dynamic systems [...]. The authors of these studies have not been interested in the problem of the foundations of statistical mechanics which is our primary interest in this book. Their aim was to obtain the results in the most general form; in particular all these results pertain equally to the systems with only a few degrees of freedom as well as to the systems with a very large number of degrees of freedom.

From our point of view we must deviate from this tendency. We would unnecessarily restrict ourselves by neglecting the special properties of the systems considered in statistical mechanics (first of all their fundamental property of having a very large number of degrees of freedom) [...]. Furthermore, we do not have any basis for demanding the possibility of substituting phase averages for the time averages of all functions; in fact the functions for which such substitution is desirable have many specific properties which make such a substitution apparent in these cases (Khinchin, 1949, p. 62).

Thus, partly in order to supplement, partly in competition to ergodic theory, Khinchin explored an approach to the ergodic problem that takes the large number of degrees of freedom as an essential ingredient, but only works for a specific class of functions, the so-called *sum functions*.

In particular, consider a Hamiltonian dynamical system $\langle \Gamma, \mathcal{A}, T, \mu \rangle$ of N point particles. That is, we assume: $x = (\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N) \in \Gamma \subset \mathbb{R}^{6N}$. A function f on Γ is a sum function if

$$(114) \quad f(x) = \sum_{i=1}^N \phi_i(x_i)$$

where $x_i = (\vec{p}_i, \vec{q}_i)$ is the molecular state of particle i .⁶¹ Under the further assumption that the Hamiltonian itself is a sum function, Khinchin proved:

KHINCHIN'S ERGODIC THEOREM: For all sum functions f there are positive constants κ_1, κ_2 such that, for all N :

$$(115) \quad \mu \left(\left\{ x \in \Gamma : \left| \frac{f(x) - \langle f \rangle_\mu}{\langle f \rangle_\mu} \right| \geq \kappa_1 N^{-1/4} \right\} \right) \leq \kappa_2 N^{-1/4}$$

In words: as N becomes larger and larger, the measure of the set where \bar{f} and $\langle f \rangle$ deviate more than a small amount goes to zero.

This theorem, then, provides an alternative strategy to address the ergodic problem: it says that time average and microcanonical phase average of sum functions will be roughly equal, at least in a very large subset of the energy hypersurface,

⁶¹Note that Khinchin does not demand that sum functions are symmetric under permutation of the particles.

provided that the number of particles is large enough. Of course, this ‘rough equality’ is much weaker than the strict equality ‘almost everywhere’ stated in the von Neumann–Birkhoff ergodic theorem. Moreover, it holds only for the sum functions (114). However, the assumption of metric transitivity is not needed here; nor is any of the more stringent properties of the ergodic hierarchy.

The advantages of this approach to the ergodic problem are clear: first, one avoids the problem that ergodic properties are hard to come by for physically interesting systems. Second, an important role is allotted to the large number of degrees of freedom, which, as noted above, seems a necessary, or at least welcome ingredient in any explanation of thermal behaviour,⁶² and thirdly a physically motivated choice for special functions has been made.

However, there are also problems and drawbacks. First, with regard to the “infinite-times” problem (cf. p. 1010), Khinchin’s approach fares no better or worse than the original ergodic approach. Second, since the rough equality does not hold “almost everywhere” but outside of a subset whose measure becomes small when N is large, the measure-zero problem of ergodic theory (p. 1010) is now replaced by a so-called “measure-epsilon problem”: if we wish to conclude that in practice the time average and the phase average are (roughly) equal, we should argue that the set for which this does not hold, i.e. the set in the left-hand side of (115) is negligible. This problem is worse than the 0 measure-zero problem. For example, we cannot argue that ensembles whose density functions have support in such sets are excluded by an appeal to absolute continuity or translation continuity (cf. the discussion on p. 1003). Further, if we wish to apply the result to systems that are indeed not metrically transitive, there may be integrals of the equations of motion that lock the trajectory of the system into a tiny subset of Γ for all times, in which case such a set cannot be neglected for practical purposes (cf. [Farquhar, 1964]).

Khinchin argued that the majority of physically important phase functions that one encounters in statistical mechanics are sum functions (cf. [Khinchin, 1949, p. 63,97]). However, this view is clearly too narrow from a physical point of view. It means that all quantities that depend on correlations or interactions between the particles are excluded.

Finally there is the ‘methodological paradox’ [Khinchin, 1949, p. 41–43]. This refers to the fact that Khinchin had to assume that the Hamiltonian itself is also a sum function. Let me emphasize that this assumption is *not* made just for the purpose of letting the Hamiltonian be one of the functions to which the theorem applies; the assumption is crucial to the very derivation of the theorem. As Khinchin clearly notes, this is paradoxical because for an equilibrium state to arise at all, it is essential that the particles can interact (e.g. collide), while this possibility is denied when the Hamiltonian is a sum function.

In Khinchin’s view, the assumption should therefore not be taken literally. In-

⁶²The point can be debated, of course. Some authors argue that small systems can show thermal behaviour too, which statistical mechanics then ought explain. However, the very definition of thermal quantities (like temperature etc.) for such small systems is more controversial [Hill, 1987; Feshbach, 1987; Rugh, 2001; Gross and Votyakov, 2000].

stead, one should assume that there really are interaction terms in the Hamiltonian, but that they manifest themselves only at short distances between the particles, so that they can be neglected, except on a tiny part of phase space. Still, it remains a curious feature of his work that his theorem is intended to apply in situations that are inconsistent with the very assumptions needed to derive it (cf. [Morrison, 2000, p. 46-47]). As we shall see in the next paragraph, later work has removed this paradox, as well as many other shortcomings of Khinchin's approach.

The theory of the thermodynamic limit

The approach initiated by Khinchin has been taken further by van der Linde and Mazur (1963), and merged with independent work of van Hove, Yang and Lee, Fisher, Griffiths, Minlos, Ruelle, Lanford and others, to develop, in the late 60s and early 70s, into what is sometimes called the 'rigorous results' approach or the 'theory of the thermodynamic limit'. The most useful references are [Ruelle, 1969; Lanford, 1973; Martin-Löf, 1979]. The following is primarily based on [Lanford, 1973], which is the most accessible and also the most relevant for our purposes, since it explicitly addresses the ergodic problem, and on [van Lith, 2001b].

As in Khinchin's work, this approach aims to provide an explanatory programme for the thermal behaviour of macroscopic bodies in equilibrium by relying mostly on the following central points,

- One adopts the microcanonical measure on phase space.
- the observable quantities are phase functions F of a special kind (see below).
- The number of particles N is extremely large.

It is shown that, under some conditions, in the 'thermodynamic limit', to be specified below, the microcanonical probability distribution for F/N becomes concentrated within an narrow region around some fixed value. This result is similar to Khinchin's ergodic theorem. However, as we shall see, the present result is more powerful, while the assumptions needed are much weaker.

To start of, we assume a Hamiltonian, of the form

$$(116) \quad H(x) = \sum_i^N \frac{\vec{p}_i^2}{2m} + U(\vec{q}_1, \dots, \vec{q}_N).$$

defined on the phase space Γ for N particles. For technical reasons, it is more convenient and simpler to work in the configuration space, and ignore the momenta. Consider a sequence of functions $F(\vec{q}_1, \dots, \vec{q}_n)$, $n = 1, 2, \dots$ with an indefinite number of arguments, or, what amounts to the same thing, a single function F defined on

$$(117) \quad \cup_{n=1}^{\infty} (\mathbb{R}^3)^n.$$

Such a function is called an ‘observable’ if it possesses the following properties:

- (a). Continuity: For each n , $F(\vec{q}_1, \dots, \vec{q}_n)$ is a continuous function on \mathbb{R}^{3n}
- (b). Symmetry: For each n , $F(\vec{q}_1, \dots, \vec{q}_n)$ is invariant under permutation of its arguments.
- (c). Translation invariance: For each n , and each $\vec{a} \in \mathbb{R}^3$, $F(\vec{q}_1 + \vec{a}, \dots, \vec{q}_n + \vec{a}) = F(\vec{q}_1, \dots, \vec{q}_n)$
- (d). Normalization: $F(\vec{q}_1) = 0$
- (e). Finite range: There exists a real number $R \in \mathbb{R}$ such that, for each n , the following holds: Suppose we divide the n particles into two clusters labeled by $i = 1, \dots, m$, and $i' = 1, \dots, m'$, where $m + m' = n$. If $|\vec{q}_i - \vec{q}_{i'}| > R$ for all i, i' , then $F(\vec{q}_1, \dots, \vec{q}_m; \vec{q}_1, \dots, \vec{q}_{m'}) = F(\vec{q}_1, \dots, \vec{q}_m) + F(\vec{q}_1, \dots, \vec{q}_{m'})$.

For the most part, these conditions are natural and self-explanatory. Note that the symmetry condition (b) is very powerful. It may be compared to Boltzmann’s (1877b) combinatorial approach in which it was argued that macrostates occupy an overwhelmingly large part of phases space due to their invariance under permutations of the particles (see §4.4). Note further that condition (e) implies that F reduces to a sum function if all particles are sufficiently far from each other. It also means that the observables characterized by Lanford may be expected to correspond to extensive quantities only. (Recall that a thermodynamical quantity is called extensive if it scales proportionally to the size of the system, and intensive if it remains independent of the system size.) In the present approach, intensive quantities (like temperature and pressure) are thus not represented as observables, but rather identified with appropriate derivatives of other quantities, after we have passed to the thermodynamical limit.

Further, it is assumed that the potential energy function U in (116) also satisfies the above conditions. In addition, the potential energy is assumed to be *stable*,⁶³ i.e.:

- (f). Stability: There is a number $B \in \mathbb{R}$, such that, for all n and all $\vec{q}_1, \dots, \vec{q}_n$:

$$(118) \quad U(\vec{q}_1, \dots, \vec{q}_n) \geq -nB.$$

This condition — which would be violated e.g. for Newtonian gravitational interaction — avoids that as n becomes large, the potential energy per particle goes to minus infinity, i.e., it avoids a collapse of the system.

For some results it is useful to impose an even stronger condition:

⁶³Strictly speaking, condition (f) is not needed for the existence of the thermodynamic limit for the configurational microcanonical measure. It is needed, however, when these results are extended to phase space (or when using the canonical measure). Note also that the term ‘stability’ here refers to an extensive lower bound of the Hamiltonian. This should be distinguished from thermodynamic concept of stability, which is expressed by the concavity of the entropy function (cf. p. 940).

(f'.) Superstability: The potential energy U is called superstable if, for every continuous function Φ of compact support in \mathbb{R}^3 :

$$(119) \quad U(\vec{q}_1, \dots, \vec{q}_N) + \lambda \sum_{i \neq j} \Phi(\vec{q}_i - \vec{q}_j)$$

is stable for a sufficiently small choice of $\lambda > 0$. In other words, a stable potential is superstable if it remains stable when perturbed slightly by a continuous finite-range two-body interaction potential.

As in Khinchin's approach, the assumption (f) or (f') is not just needed because one would like to count the potential energy among the class of observables; rather it is crucial to the proof of the existence of the thermodynamic limit. Of course, the assumption that the interaction potential is continuous and of finite range is still too restrictive to model realistic inter-molecular forces. As Lanford notes, one can weaken condition (e) to a condition of 'weakly tempered' potentials,⁶⁴ dropping off quickly with distance (cf. [Fisher, 1964, p. 386; Ruelle, 1969, p. 32], although this complicates the technical details of the proofs. Again, it is clear, however, that some such condition on temperedness of the long range interactions is needed, if only to avoid another catastrophe, namely that the potential energy per particle goes to $+\infty$ as n increases, so that system might tend to explode. (As could happen, e.g. for a system of charges interacting by purely repulsive Coulomb forces.)

Now, with the assumptions in place, the idea is as follows. Choose a given potential U and an observable F obeying the above conditions. Pick two numbers u and ρ , that will respectively represent the (potential) energy per particle and the particle density (in the limit as N gets large), a bounded open region $\Lambda \subset \mathbb{R}^3$, and a large integer N , such that $\frac{N}{V(\Lambda)} \approx \rho$. (Here, $V(\Lambda)$ denotes the volume of Λ .) Further, choose a small number $\delta u > 0$, and construct the (thickened) energy hypersurface in configuration space, i.e. the shell:

$$(120) \quad \Omega_{\Lambda, N, u, \delta u} = \left\{ (\vec{q}_1, \dots, \vec{q}_N) \in \Lambda^N : \frac{U(\vec{q}_1, \dots, \vec{q}_N)}{N} \in (u - \delta u, u + \delta u) \right\}.$$

Let μ denote the Lebesgue measure on Λ^N ; its (normalized) restriction to the above set may then be called the 'thickened configurational microcanonical measure'. Note that

$$(121) \quad \omega^{\text{cf}}(E) := \int_{\Lambda^N} d\vec{q}_1 \cdots d\vec{q}_N \delta(U(\vec{q}_1 \dots \vec{q}_N) - E)$$

may be considered as the configurational analogue of the structure function (41). Thus

⁶⁴If, for simplicity, the potential U is a sum of pair interactions $U = \sum_{i \neq j} \phi(\vec{q}_i - \vec{q}_j)$, it is weakly tempered iff there are real constants $R, D, \epsilon > 0$, such that $\phi(\vec{r}) \leq D \|\vec{r}\|^{3+\epsilon}$ when $\|\vec{r}\| \geq R$.

$$(122) \quad \mu(\Omega_{\Lambda, N, u, \delta u}) = \int_{\Lambda^N} d\vec{q}_1 \cdots d\vec{q}_N \mathbf{1}_{(u-\delta u, u+\delta u)}(U/N) = \int_{N(u-\delta u)}^{N(u+\delta u)} dE \omega^{\text{cf}}(E),$$

so that $\frac{1}{2N\delta u} \mu(\Omega_{\Lambda, N, u, \delta u})$ provides a thickened or smoothed version of this configurational structure function. The reason for working with this thickened hypershell instead of the thin hypersurface is of course to circumvent singularities that may appear in the latter. In any case, we may anticipate that, when δu is small, this expression will represent the configurational part of the microcanonical entropy (84). A further factor $1/N!$ may be added to give this entropy a chance of becoming extensive.⁶⁵ (See also paragraph §5.2.

We are interested in the probability distribution of F/N with respect to this thickened microcanonical measure on configuration space. For this purpose, pick an arbitrary open interval J , and define

$$(123) \quad \mathcal{V}(\Lambda, N, u, \delta u, F, J) := \frac{1}{N!} \mu \left(\left\{ (\vec{q}_1, \dots, \vec{q}_N) \in \Omega_{u, \delta u} : \frac{F(\vec{q}_1, \dots, \vec{q}_N)}{N} \in J \right\} \right).$$

So,

$$(124) \quad \frac{1}{\mu(\Omega_{\Lambda, N, u, \delta u})} \mathcal{V}(\Lambda, N, u, \delta u, F, J) = \frac{\mathcal{V}(\Lambda, N, u, \delta u, F, J)}{\mathcal{V}(\Lambda, N, u, \delta u, F, \mathbb{R})}$$

gives the probability that F/N lies in the interval J with respect to the above microcanonical measure.

We wish to study the behaviour of this probability in the thermodynamic limit, i.e. as N becomes large, and $V(\Lambda)$ grows proportional to N , such that $N/V(\Lambda) = \rho$. This behaviour will depend on the precise details of the limiting procedure, in particular on the shape of Λ . Lanford chooses to take the limit in the sense of van Hove: A sequence of bounded open regions Λ in \mathbb{R}^3 is said to become infinitely large in the sense of Van Hove if, for all $r > 0$, the volume of the set of all points within a distance r from the boundary of Λ , divided by the volume of Λ , goes to zero as N goes to infinity. In other words, the volume of points close to the surface becomes negligible compared to the volume of the interior. This avoids that surface effects could play a role in the limiting behaviour — and eliminates the worry that interactions with the walls of the container should have been taken into account.

Now, the first major result is:

(EXISTENCE OF THE THERMODYNAMIC LIMIT.) As $N \rightarrow \infty$, and Λ becomes infinitely large in the sense of Van Hove, in such a way that $N/V(\Lambda) = \rho$, then either of the following cases holds:

(α). $\mathcal{V}(\Lambda, N, u, \delta u, F, J)$ goes to zero faster than exponentially in N ,
or:

⁶⁵For example, if the system is an ideal gas, i.e. if $U(\vec{q}_1, \dots, \vec{q}_N) \equiv 0$, one will have $\omega^{\text{cf}}(E) = V^N = \left(\frac{N}{\rho}\right)^N$, so that $\ln \frac{1}{N!} \omega^{\text{cf}}(E)$ scales proportionally to N , but $\ln \omega^{\text{cf}}(E)$ does not.

(β). $\mathcal{V}(\Lambda, N, u, \delta u, F, J) \approx e^{Ns(\rho, u, \delta u, F, J)}$ where $s(\rho, F, J)$ does not depend on Λ or N , except through the ration $\frac{N}{V(\Lambda)} = \rho$.

In other words, this result asserts the existence of

$$(125) \quad s(\rho, u, \delta u, F, J) := \lim_{N \rightarrow \infty} \frac{1}{N} \ln \mathcal{V}(\Lambda, N, u, \delta u, F, J)$$

where s is either finite or $-\infty$. (The possibility that $s = -\infty$ for *all* values of the arguments of s is further ignored.) This already gives some clue for how the probability (123) behaves as a function of J . If J_1 and J_2 are two open intervals, N is large, and we suppress the other variables for notational convenience, we expect:

$$(126) \quad \frac{\mu(\frac{F}{N} \in J_1)}{\mu(\frac{F}{N} \in J_2)} = \frac{\mathcal{V}(J_1)}{\mathcal{V}(J_2)} \approx e^{N(s(J_1) - s(J_2))}.$$

If $s(J_2) > s(J_1)$, this ratio goes to zero exponentially in N . Thus, for large systems, the probability $\mu(\frac{F}{N} \in J)$ will only be appreciable for those open intervals J for which $s(J)$ is large.

A stronger statement can be obtained as follows. Associated with the set function $s(J)$ one may define a point function s :

$$(127) \quad s(x) := \inf_{\substack{J \ni x \\ J \text{ open}}} s(J)$$

It can then be shown that, conversely, for all open J :

$$(128) \quad s(J) = \sup_{x \in J} s(x)$$

Moreover, — and this is the second major result — one can show:

$$(129) \quad s(x) \text{ is concave.}$$

Further, $s(x)$ is finite on an open convex subset of its domain [Lanford, 1973, p. 26].

Now, it is evident that a concave function $s(x)$ may have three general shapes: It either achieves its maximum value: (i) never; (ii) exactly once, say in some point x_0 ; or (iii) on some interval. In case (i), F/N ‘escapes to infinity’ in the thermodynamic limit; this case can be excluded by imposing the superstability condition (f'). Case (ii) is, for our purpose, the most interesting one. In this case, we may consider intervals $J_2 = (x_0 - \epsilon, x_0 + \epsilon)$, for arbitrarily small $\epsilon > 0$ and J_1 any open interval that does not contain x_0 ; infer from (127,128) that $s(J_2) > s(J_1)$, and conclude from (126) that the relative probability for F/N to take a value in J_2 rather than J_1 goes to zero exponentially with the size of the system.

Thus we get the desired result: As N becomes larger and larger, the probability distribution of F/N approaches a delta function. Or in other words, the function F/N becomes roughly constant on an overwhelmingly large part of the configurational energy-hypershell:

$$(130) \lim_{N \rightarrow \infty} \mu \left(\left\{ (\vec{q}_1, \dots, \vec{q}_N) \in \Omega_{\Lambda, N, u, \delta u} : \left| \frac{F(\vec{q}_1, \dots, \vec{q}_N)}{N} - x_0 \right| > \epsilon \right\} \right) = 0$$

In case (iii), finally, one can only conclude that the probability distribution becomes concentrated on some interval, but that its behaviour inside this interval remains undetermined. One can show, if this interval is bounded, that this case is connected to phase transitions (but see [p. 12, 58 for caveats]).⁶⁶

Remarks.

1. Phase transitions. First, it is obviously an immense merit of the theory of the thermodynamic limit that, in contrast to ergodic theory, it is, in principle, capable of explaining and predicting the occurrence of phase transitions from a model of the microscopic interaction, in further work often in conjunction with renormalization techniques. Indeed, this capability is its major claim to fame, quite apart from what it has to say about the ergodic problem. What is more, it is often argued that phase transitions are strictly impossible in any finite system, and thus absolutely *require* the thermodynamic limit [Styer, 2004; Kadanoff, 2000].

This argument raises the problem that our experience, including that of phase transitions in real physical bodies, always deals with finite systems. A theory that presents an account of phase transitions only in the thermodynamic limit, must then surely be regarded as an idealization. This conclusion will not come as a shock many physicists, since idealizations are ubiquitous in theoretical physics. Yet a curious point is that this particular idealization seems to be ‘uncontrollable’. See [Sklar, 2002] and [Liu, 1999; Callender, 2001; Batterman, 2005] for further discussion. I also note that an alternative approach has been proposed recently. In this view phase transitions are associated with topology changes in the microcanonical hypersurface $\{x : H(x) = E\}$ with varying E . The crucial distinction from the theory of the thermodynamic limit is, of course, is that such topology changes may occur in finite, — indeed even in small — systems (cf. [Gross, 1997; Gross and Votyakov, 2000; Casetti et al, 2003]). However this may be, I shall focus below on the virtues of the thermodynamic limit for the ergodic problem.

2. The ergodic problem revisited. When compared to ergodic theory or Khinchin’s approach, the theory of the thermodynamic limit has much to speak in

⁶⁶To see the connection (loosely), note that if one removes the condition $F/N \in J$ from the definition (123) — or equivalently, chooses $J = \mathbb{R}$ —, then s in (125) can be interpreted as the (thickened, configurational) microcanonical entropy per particle. Considered now as a function of the open interval $(u - \delta u, u + \delta u)$, s has the same properties as established for $s(J)$, since U itself belongs to the class of observables. Thus, here too, there exists a point function $s(u)$ analogous to (127), and this function is concave (Actually, if we restore one more variable in the notation, and write $s(\rho, u)$, the function is concave in both variables). In case (iii), therefore, this function is constant in u over some interval, say $[u'_0, u''_0]$. This means that there is then a range of thermodynamical states with the same temperature $T = (\frac{\partial s}{\partial u})_\rho^{-1}$, for a range of values of u and ρ , which is just what happens in the condensation phase transition in a van der Waals gas.

its favour. As in Khinchin's work, the problem of establishing metric transitivity for physically interesting systems does not arise, because the approach does not need to assume it. Further, as in Khinchin's work, the approach works only for special functions. But the class of functions singled out by the assumptions (a-f.) or (a-f'.) above is not restricted to (symmetric) sum functions, and allows for short-range interactions between the particles. Thus, unlike Khinchin, there is no methodological paradox (cf. p.1019).

Yet one might still question whether these assumptions are not too restrictive for physically interesting systems. On the one hand, it is clear that some conditions on temperedness and stability are needed to rule out catastrophic behaviour in the thermodynamic limit, such as implosion or explosion of the system. On the other hand, these assumptions are still too strong to model realistic thermal systems. The Coulomb interaction, which according to [Lieb and Lebowitz, 1973, p. 138] is "the true potential relevant for real matter", is neither tempered nor stable. A tour de force, performed by Lenard, Dyson, Lebowitz and Lieb, has been to extend the main results of the theory of the thermodynamical limit to systems interacting purely by Coulomb forces (if the net charge of the system is zero or small), both classically and quantum mechanically (for fermions) (see [Lieb, 1976, and literature cited therein]). This result, then, should cover most microscopic models of ordinary matter, as long as relativistic effects and magnetic forces can be ignored. But note that this extension is obtained by use of the canonical, rather than the microcanonical measure, and in view of the examples of non-equivalence of these ensembles (cf. p. 5.5) one might worry whether this result applies to ordinary matter in metastable states (like supersaturated vapours, and superheated or supercooled liquids).

Another remarkable point is that, unlike Khinchin's result (115), the result (130) does not refer to time averages at all. Instead, the *instantaneous* value of F/N is found to be almost constant for a large subset of the configurational energy hypersurface. Hence, there is also no problem with the infinite time limit (cf. p. 1010. Indeed, dynamics or time evolutions play no role whatsoever in the present results, and the contrast to the programme of ergodic theory is accordingly much more pronounced than in Khinchin's approach.

3. Problems left. What is left, in comparison to those two approaches to the ergodic problem, are two problems. First, there is still the question of how to motivate the choice for the configurational microcanonical measure (i.e. the normalized Lebesgue measure restricted to the energy hypershell). Lanford is explicit that the theory of the thermodynamic limit offers no help in this question:

It is a much more profound problem to understand why events which are very improbable with respect to Lebesgue measure do not occur in nature. I, unfortunately, have nothing to say about this latter problem. [Lanford, 1973, p. 2].

For this purpose, one would thus have to fall back on other attempts at motivation (cf. p. 1003).

Secondly, there is the measure-epsilon problem (cf. p. 1019). The desired equality $F/N \approx x$ holds, according to (130), if N is large, outside of a set of small measure. Can we conclude that this set is negligible, or that its states do not occur in nature? In fact, the result (130) instantaneous values is so strong that one ought to be careful of not claiming too much. For example, it would be wrong to claim that for macroscopical systems (i.e. with $N \approx 10^{27}$), the set in the left-hand side of (130) does not occur in nature. Instead, it remains a brute fact of experience that macroscopic systems also occur in non-equilibrium states. In such states, observable quantities take instantaneous values that vary appreciably over time, and thus differ from their microcanonical average. Therefore, their microstate must then be located inside the set of tiny measure that one would like to neglect. Of course, one might argue differently if N is larger still, say $N = 10^{100}$ but this only illustrates the ‘uncontrollability’ of the idealization involved in this limit, i.e. one still lacks control over how large N must be to be sure that the thermodynamic limit is a reasonable substitute for a finite system.

Further points. Other points, having no counterpart in the approaches discussed previously, are the following. The approach hinges on a very delicately construed sequence of limits. We first have to take the thickened energy shell, then take N, Λ to infinity in the sense of van Hove, finally take δu to zero. But one may ask whether this is clearly and obviously the right thing to do, since there are alternative and non-equivalent limits (the sense of Fisher), the order of the limits clearly do not commute (the thickness of the energy hypershell is proportional to $N\delta u$), and other procedures like the ‘continuum limit’ [Compagner, 1989] have also been proposed.

Finally, in order to make full contact to classical statistical mechanics, one still has to lift restriction to configuration space, and work on phase space. [Lanford, 1973, p. 2] leaves this as a “straightforward exercise” to the reader. Let’s see if we can fill in the details.

Suppose we start from a thickened microcanonical measure on phase space, with the same thickness $2N\delta u$, around a total energy value of $E_0 = Ne_0$. Its probability density is then given by

$$(131) \quad \rho_{Ne_0, N\delta u}(\vec{p}_1, \dots, \vec{p}_N; \vec{q}_1 \dots \vec{q}_N) = \frac{1}{2N\delta u} \int_{E_0 - N\delta u}^{E_0 + N\delta u} \frac{1}{\omega(E)} \delta(H(x) - E) dE$$

For the Hamiltonian (116), the integral over the momenta can be performed (as was shown by Boltzmann [1868] (cf. Eqn (43)). This yields a marginal density

$$(132) \quad \bar{\rho}_{Ne_0, N\delta u}(\vec{q}_1, \dots, \vec{q}_N) = \frac{1}{2N\delta u} \frac{2m\pi^{3N/2}}{\Gamma(\frac{3N}{2})} \int_{E_0 - N\delta u}^{E_0 + N\delta u} \frac{1}{\omega(E)} (2m(E - U(q)))^{(3N-2)/2} dE$$

This is not quite the normalized Lebesgue measure on configuration space employed by Lanford, but since the factor $(2m(E - U(q)))^{(3N-2)/2}$ is a continuous

function of U , — at least if $E_0 - N\delta u - U > 0$ — it is absolutely continuous with respect to the Lebesgue measure on the shell, and will converge to it in the limit $\delta u \rightarrow 0$.

But in a full phase space setting, the physical quantities can also depend on the momenta, i.e., they will be functions $F(\vec{p}_1, \dots, \vec{p}_N; \vec{q}_1 \dots \vec{q}_N)$ and, even if one assumes the same conditions (a-f) as before for their dependence on the second group of arguments, their probability distribution cannot always be determined from the configurational microcanonical measure. For example, let F_1 and F_2 be two observables on configuration space, for which F_1/N and F_2/N converge to different values in the thermodynamical limit, say x_1 and x_2 , and let G be any symmetric function of the momenta that takes two different values each with probability 1/2. For example, take

$$(133) \quad G(\vec{p}_1, \dots, \vec{p}_N) = \begin{cases} 1 & \text{if } \sum_i \vec{p}_i \cdot \vec{n} \geq 0, \\ 0 & \text{elsewhere.} \end{cases},$$

for some fixed unit vector \vec{n} . Now consider the following function on phase space:

$$(134) \quad A(\vec{p}_1, \dots, \vec{p}_N; \vec{q}_1 \dots \vec{q}_N) = G(\vec{p}_1, \dots, \vec{p}_N) F_1(\vec{q}_1 \dots \vec{q}_N) + G'(\vec{p}_1, \dots, \vec{p}_N) F_2(\vec{q}_1 \dots \vec{q}_N),$$

where $G' = 1 - G$. If we first integrate over the momenta, we obtain $\tilde{A} = \frac{1}{2}(F_1 + F_2)$, which converges in the thermodynamical limit to $\frac{1}{2}(x_1 + x_2)$. However, it would be wrong to conclude that A is nearly equal to $\frac{1}{2}(x_1 + x_2)$ ($(x_1 + x_2)/2$) in an overwhelmingly large part of phase space. Instead, it is nearly equal to x_1 on (roughly) half the available phase space and nearly equal to x_2 on the remaining half.

The extension of (130) to phase space functions will thus demand extra assumptions on the form of such functions; for example, that their dependence on the momenta comes only as some function of the kinetic energy, i.e.

$$(135) \quad A(\vec{p}_1, \dots, \vec{p}_N; \vec{q}_1, \dots, \vec{q}_N) = \psi\left(\sum \frac{\vec{p}_i^2}{2m}\right) + F(\vec{q}_1, \dots, \vec{q}_N)$$

for some continuous function ψ .

6.4 Lanford's approach to the Boltzmann equation

We now turn to consider some modern approaches to non-equilibrium statistical mechanics. Of these, the approach developed by Lanford and others (cf. [Lanford, 1975; Lanford, 1976; Lanford, 1981; Spohn, 1991; Cercignani *et al.*, 1994]) deserves special attention because it stays conceptually closer to Boltzmann's 1872 work on the Boltzmann equation and the H -theorem than any other modern approach to statistical physics. Also, the problem Lanford raised and tried to answer is one of no less importance than the famous reversibility and recurrence objections. Furthermore, the results obtained are the best efforts so far to show that a statistical reading of the Boltzmann equation or the H -theorem might hold for the hard spheres gas.

The question Lanford raised is that of the consistency of the Boltzmann equation and the underlying Hamiltonian dynamics. Indeed, if we consider the microstate of a mechanical system such as a dilute gas, it seems we can provide two competing accounts of its time evolution.

(1) On the one hand, given the mechanical microstate x_0 of a gas, we can form the distribution of state $f(\vec{r}, \vec{v})$, such that $f(\vec{r}, \vec{v})d^3\vec{v}d^3\vec{r}$ gives the relative number of molecules with a position between \vec{r} and $\vec{r} + d^3\vec{r}$ and velocity between \vec{v} and $\vec{v} + d^3\vec{v}$. Presumably, this distribution should be uniquely determined by the microstate x_0 . Let us make this dependence explicit by adopting the notation $f^{[x_0]}$. This function, then, should ideally serve as an initial condition for the Boltzmann equation (48), and solving this equation — assuming, that is, that it, that it has a unique solution — would give us the shape of the distribution function at a later time, $f_t^{[x_0]}(\vec{r}, \vec{v})$.

(2) On the other hand, we can evolve the microstate x_0 for a time t with the help of the Hamiltonian equations. That will give us $x_t = T_t x_0$. This later state x_t will then also determine a distribution of state $f^{[x_t]}(\vec{r}, \vec{v})$.

It is a sensible question whether these two ways of obtaining a later distribution of state from an initial microstate are the same, i.e. whether the two time evolutions are consistent. In other words, the problem is whether the diagram below commutes:

$$(136) \quad \begin{array}{ccc} x_0 & \xrightarrow{\text{Hamilton}} & x_t \\ \downarrow & & \downarrow \\ f^{[x_0]} & \xrightarrow{\text{Boltzmann}} & f_t^{[x_0]} \stackrel{?}{=} f^{[x_t]} \end{array}$$

The first issue that has to be resolved here is the precise relation between a microstate and the distribution of state f . It is obvious that, in so far as this function represents the physical property of a gas system, it should be determined by the momentary microstate x . It is also clear, that in so far as it is assumed to be continuous and differentiable in time in order to obey the Boltzmann equation, this cannot be literally and exactly true.

So let us assume, as Boltzmann did, that the gas consists of N hard spheres, each of diameter d and mass m , contained in some fixed bounded spatial region Λ with volume $|\Lambda| = V$. Given a microstate x of the system one can form the ‘exact’ distribution of state:

$$(137) \quad F^{[x]}(\vec{r}, \vec{v}) := \frac{1}{N} \sum_i^N \delta^3(\vec{r} - \vec{q}_i) \delta^3(\vec{v} - \frac{\vec{p}_i}{m}).$$

This distribution is, of course, not a proper function, and being non-continuous and non-differentiable, clearly not a suitable object to plug into the Boltzmann equation. However, one may reasonably suppose that one ought to be able to express Boltzmann’s ideas in a limit in which the number of particles, N , goes to infinity. However, this limit clearly must be executed with care.

On the one hand, one ought to keep the gas dilute, so that collisions involving three or more particles will be rare enough so that they can safely be ignored in comparison to two-particle collisions. On the other hand, the gas must not be so dilute that collisions are altogether too rare to contribute to a change of f . The appropriate limit to consider, as Lanford argues, is the so-called Boltzmann-Grad limit in which $N \rightarrow \infty$, and:⁶⁷

$$(138) \quad \frac{Nd^2}{V} = \text{constant} > 0.$$

Denote this limit as “ $N \xrightarrow{BG} \infty$ ”, where it is implicitly understood that $d \propto N^{-1/2}$. The hope is then that in this Boltzmann-Grad limit, the exact distribution $F^{[x^N]}$ will tend to a continuous function that can be taken as an appropriate initial condition for the Boltzmann equation. For this purpose, one has to introduce a relevant notion of convergence for distributions on the μ -space $\Lambda \times \mathbb{R}^3$. A reasonable choice is to say that an arbitrary sequence of distributions f_n (either proper density functions or in the distributional sense) converges to a distribution f , $f_n \rightarrow f$, iff the following conditions hold:

For each rectangular parallelepiped $\Delta \subset \Lambda \times \mathbb{R}^3$:

$$(139) \quad \lim_{n \rightarrow \infty} \int_{\Delta} f_n d^3\vec{r} d^3\vec{v} = \int_{\Delta} f d^3\vec{r} d^3\vec{v},$$

$$(140) \quad \text{and} \quad \lim_{n \rightarrow \infty} \int \vec{v}^2 f_n d^3\vec{r} d^3\vec{v} = \int \vec{v}^2 f d^3\vec{r} d^3\vec{v},$$

where the second condition is meant to guarantee the convergence of the mean kinetic energy.

It is also convenient to introduce some distance function between (proper or improper) distributions that quantifies the sense in which one distribution is close to another in the above sense. That is to say, one might define some distance $d(f, g)$ between density functions on $\Lambda \times \mathbb{R}^3$ such that

$$(141) \quad d(f_n, f) \rightarrow 0 \implies f_n \rightarrow f.$$

There are many distance functions that could do this job, but I won't go into the question of how to pick out a particular one.

The hope is then, to repeat, that $F^{[x^N]} \rightarrow f$ in the above sense when $N \xrightarrow{BG} \infty$, where f is sufficiently smooth to serve as an initial condition in the Boltzmann equation, and that with this definition, the Boltzmannian and Hamiltonian evolution become consistent in the sense that the diagram (136) commutes. But clearly this will still be a delicate matter. Indeed, increasing N means a transition from one mechanical system to another with more particles. But there is no obvious

⁶⁷The condition can be explained by the hand-waving argument that Nd^2/V is proportional to the ‘mean free path’, i.e. a typical scale for the distance traveled by a particle between collisions, or also by noting that the collision integral in the Boltzmann equation is proportional to Nd^2/V , so that by keeping this combination constant, we keep the Boltzmann equation unchanged.

algorithm to construct the state x^{N+1} from x^N , and thus no way to enforce convergence on the level of individual states.

Still, one might entertain an optimistic guess, which, if true, would solve the consistency problem between the Boltzmann and the Hamiltonian evolution in an approximate fashion if N is very large.

OPTIMISTIC GUESS: If $F^{[x_0^N]}$ is near to f then $F^{[x_t^N]}$ is near to f_t for all $t > 0$, and where f_t is the solution of the Boltzmann equation with initial condition f .

As [Lanford, 1976] points out, the optimistic guess cannot be right. This is an immediate consequence of the reversibility objection: Indeed, suppose it were true for all $x \in \Gamma$, and $t > 0$. (Here, we momentarily drop the superscript N from x^N to relieve the notation.) Consider the phase point Rx obtained from x by reversing all momenta: $R(\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N) = (\vec{q}_1, -\vec{p}_1; \dots; \vec{q}_N, -\vec{p}_N)$. If $F^{[x]}(\vec{r}, \vec{v})$ is near to some distribution $f(\vec{r}, \vec{v})$, then $F^{[Rx]}(\vec{r}, \vec{v})$ is near to $f(\vec{r}, -\vec{v})$. But as x evolves to x_t , Rx_t evolves to $T_t Rx_t = RT_{-t}x_t = Rx$. Hence $F^{[T_t Rx_t]}(\vec{r}, \vec{v}) = F^{[Rx]}(\vec{r}, \vec{v})$ is near to $f(\vec{r}, -\vec{v})$. But the validity of the conjecture for Rx_t would require that $F^{[T_t Rx_t]}(\vec{r}, \vec{v})$ is near to $f_t(\vec{r}, -\vec{v})$ and these two distributions of state are definitely not near to each other, except in some trivial cases.

But even though the optimistic guess is false in general, one might hope that it is ‘very likely’ to be true, with some overwhelming probability, at least for some finite stretch of time. In order to make such a strategy more explicit, Lanford takes recourse to a probability measure on Γ , or more precisely a sequence of probability measures on the sequence of Γ_N ’s.

Apart from thus introducing a statistical element into what otherwise would have remained a purely kinetic theory account of the problem, there is a definite advantage to this procedure. As mentioned above, there is no obvious algorithm to construct a sequence of microstates in the Boltzmann-Grad limit. But for measures this is different. The microcanonical measure, for example is not just a measure for the energy hypersurface of one N -particles-system; it defines an algorithmic sequence of such measures for each N .

In the light of this discussion, we can now state Lanford’s theorem as follows [Lanford, 1975; 1976]:

LANFORD’S THEOREM: Let $t \mapsto f_t$ be some solution of the Boltzmann equation, say for $t \in [0, a) \subset \mathbb{R}$. For each N , let Δ_N denote the set in the phase space Γ_N of N particles, on which $F^{[x^N]}$ is near to f_0 (the initial condition in the solution of the Boltzmann equation) in the sense that for some chosen distance function d and for tolerance $\epsilon > 0$:

$$(142) \quad \Delta_N = \{x^N \in \Gamma_N : d(F^{[x^N]}, f_0) < \epsilon\}.$$

Further, for each N , conditionalize the microcanonical measure μ_N on Δ_N :

$$(143) \quad \mu_{\Delta,N}(\cdot) := \mu_N(\cdot|\Delta_N).$$

In other words, $\mu_{\Delta,N}$ is a sequence of measures on the various Γ_N that assign measure 1 to the set of microstates $x^N \in \Gamma_N$ that are close to f_0 in the sense that $d(F^{[x^N]}, f_0) < \epsilon$.

Then: $\exists \tau, 0 < \tau < a$ such that for all t with $0 < t < \tau$:

$$(144) \quad \mu_{\Delta,N}(\{x^N \in \Gamma_N : d(F^{[x_t^N]}, f_t) < \epsilon\}) > 1 - \delta$$

where $\delta \rightarrow 0$ as both $\epsilon \rightarrow 0$ and $N \xrightarrow{BG} \infty$.

In other words: as judged from the microcanonical measure on Γ_N restricted to those states x^N that have their exact distribution of state close to a given initial function f_0 , a very large proportion $(1 - \delta)$ evolve by the Hamiltonian dynamics in such a way that their later exact distribution of state $F^{[x_t^N]}$ remains close to the function f_t , as evolved from f_0 by the Boltzmann equation.

Remarks

Lanford’s theorem shows that a statistical and approximate version of the Boltzmann equation can be derived from Hamiltonian mechanics and the choice of an initial condition in the Boltzmann-Grad limit. This is a remarkable achievement, that in a sense vindicates Boltzmann’s intuitions. According to [Lanford, 1976, p. 14], the theorem says that the approximate validity of the Boltzmann equation, and hence the H -theorem, can be obtained from mechanics alone and a consideration of the initial conditions.

Still the result established has several remarkable features, all of which are already acknowledged by Lanford. First, there are some drawbacks that prevent the result from having practical impact for the project of justifying the validity of the Boltzmann equation in real-life physical applications. The density of the gas behaves like N/d^3 , and in the Boltzmann-Grad limit this goes to zero. The result thus holds for extremely rarified gases. Moreover, the length of time for which the result holds, i.e. τ , depends on the constant in (138), which also provides a rough order of magnitude for the mean free path of the gas. It turns out that, by the same order of magnitude considerations, τ is roughly two fifths of the mean duration between collisions. This is a disappointingly short period: in air at room temperature and density, τ is in the order of microseconds. Thus, the theorem does not help to justify the usual applications of the Boltzmann equation to macroscopic phenomena which demand a much longer time-scale.

Yet note that the time scale is not trivially short. It would be a misunderstanding to say that the theorem establishes only the validity of the Boltzmann equation for times so short that the particles have had no chance of colliding: In two fifths of the mean duration between collisions, about 40 % of the particles have performed a collision.

Another issue is that in comparison with Boltzmann’s own derivation no explicit mention seems to have been of the *Stoßzahlansatz*. In part this is merely apparent. In a more elaborate presentation (cf. Lanford 1975, 1976), the theorem is not presented in terms of the microcanonical measure, but an arbitrary sequence of measures ν_N on (the sequence of phase spaces) Γ_N . These measures are subject to various assumptions. One is that each ν_N should be absolutely continuous with respect to the microcanonical measure μ_N , i.e. ν_N should have a proper density function

$$(145) \quad d\nu_N(x) = n_N(x_1, \dots, x_N) dx_1 \cdots dx_N$$

where $x_i = (\vec{q}_i, \vec{p}_i)$ denotes the canonical coordinates of particle i . Further, one defines, for each N and $m < N$, the reduced density functions by

$$(146) \quad n_N^{(m)}(x_1, \dots, x_m) := \frac{N!}{(N-m)! N^m} \int n_N(x_1, \dots, x_N) dx_{m+1} \cdots dx_N$$

i.e. as (slightly renormalized) marginal probability distributions for the first m particles. The crucial assumption is now that

$$(147) \quad \lim_{N \xrightarrow{\text{BG}} \infty} n_N^{(m)}(x_1, \dots, x_m) = n^{(1)}(x_1) \cdots n^{(1)}(x_m)$$

uniformly on compact subsets of $(\Lambda \times \mathbb{R}^3)^m$. This assumption (which can be shown to hold for the microcanonical measures) is easily recognized as a measure-theoretic analogy to the *Stoßzahlansatz*. It demands, in the Boltzmann-Grad limit, statistical independence of the molecular quantities for any pair or m -tuple of particles at time $t = 0$. As Lanford also makes clear, it is assumption (146) that would fail to hold if we run the construction of the reversibility objection; (i.e. if we follow the states x in Δ_N for some time t , $0t < \tau$, then reverse the momenta, and try to apply the theorem to the set $\Delta'_N = \{Rx_t : x \in \Delta_N\}$).

But another aspect is more positive. Namely: Lanford’s theorem does not need to assume explicitly that the *Stoßzahlansatz* holds *repeatedly*. Indeed a remarkable achievement is that once the factorization condition (146) holds for time $t = 0$ it will also hold for $0 < t < \tau$, albeit in a weaker form (as convergence in measure, rather than uniform convergence). This is sometimes referred to as “propagation of chaos” [Cercignani *et al.*, 1994].

But the main conceptual problem concerning Lanford’s theorem is where the apparent irreversibility or time-reversal non-invariance comes from. On this issue, various opinions have been expressed. [Lanford, 1975, p. 110] argues that irreversibility is the result of passing to the Boltzmann-Grad limit. Instead, [Lanford, 1976] argues that it is due to condition (146) plus the initial conditions (i.e.: $x_N \in \Delta_N$).

However, I would take a different position. The theorem equally holds for $-\tau < t < 0$, with the proviso that f_t is now a solution of the anti-Boltzmann equation. This means that the theorem is, in fact, invariant under time-reversal.

6.5 The BBGKY approach

The so-called BBGKY-hierarchy (named after Bogolyubov, Born, Green, Kirkwood and Yvon) is a unique amalgam of the description of Gibbs and the approach of Boltzmann. The goal of the approach is to describe the evolution of ensembles by means of reduced probability densities, and to see whether a Boltzmann-like equation can be obtained under suitable conditions — and thereby an approach to statistical equilibrium.

First, consider an arbitrary time-dependent probability density ρ_t . The evolution of ρ is determined via the Liouville-equation by the Hamiltonian:

$$(148) \quad \frac{\partial \rho_t}{\partial t} = \{H, \rho\}.$$

Central in the present approach is the observation that for relevant systems in statistical mechanics, this Hamiltonian will be symmetric under permutation of the particles. Indeed, the Hamiltonian for a system of N indistinguishable particles usually takes the form

$$(149) \quad H(\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N) = \sum_{i=1}^N \frac{\vec{p}_i^2}{2m} + \sum_i V(\vec{q}_i) + \sum_{i < j} \phi(\|q_i - q_j\|)$$

where V is the potential representing the walls of the bounded spatial region Λ , say:

$$(150) \quad V(\vec{q}) = \begin{cases} 0 & \text{if } \vec{q} \in \Lambda \\ \infty & \text{elsewhere} \end{cases}$$

and ϕ the interaction potential between particle i and j . This is not only symmetric under permutation of the particle labels, but even has the more special property that it is a sum of functions that never depend on the coordinates of more than *two* particles. (cf. the discussion in §6.3.)

Let us again use the notation $x = (\vec{q}_1, \vec{p}_1; \dots; \vec{q}_N, \vec{p}_N) = (x_1, \dots, x_N)$; with $x_i = (\vec{q}_i, \vec{p}_i)$, and consider the sequence of reduced probability density functions, defined as the marginals of ρ :

$$(151) \quad \begin{aligned} \rho^{(1)}(x_1) &:= \int \rho_t(x) dx_2 \cdots dx_N \\ &\vdots \\ \rho^{(m)}(x_1, \dots, x_m) &= \int \rho_t(x) dx_{m+1} \cdots dx_N \end{aligned}$$

Here, $\rho^{(m)}$ gives the probability density that particles $1, \dots, m$ are located at specified positions $\vec{q}_1, \dots, \vec{q}_m$ and moving with the momenta $\vec{p}_1, \dots, \vec{p}_m$, whereas all remaining particles occupy arbitrary positions and momenta.

Symmetry of the Hamiltonian need not imply symmetry of ρ . But one might argue that we may restrict ourselves to symmetric probability densities if *all* observable quantities are symmetric. In that case, it makes no observable difference when two or more particles are interchanged in the microstate and one may replace ρ by its average under all permutations without changing the expectation values of any observable quantity. However this may be, we now assume that ρ is, in fact, symmetric under permutations of the particle labels. In other words, from now on $\rho^{(m)}$ gives the probability density that any arbitrarily chosen set of m particles have the specified values for position and momentum.

The guiding idea is now that for relevant macroscopic quantities, we do not need the detailed form of the time evolution of ρ_t . Rather, it suffices to focus on no more than just a few marginals from the hierarchy (151). For example, suppose a physical quantity represented as a phase function A is a symmetric sum function on Γ :

$$(152) \quad A(x) = \sum_{i=1}^N A(x_i)$$

Then

$$(153) \quad \langle A \rangle = N \int A(x_1) \rho^{(1)}(x_1) dx_1$$

which is a considerable simplification. But this is not to say that we can compute the evolution of $\langle A \rangle$ in time so easily.

Consider in particular $\rho_t^{(1)}$ in (151). This is the *one-particle distribution function*: the probability that an arbitrary particle is in the one-particle state (\vec{p}, \vec{q}) . This distribution function is in some sense analogous to Boltzmann's f . But note: ρ_1 is a marginal probability distribution; it characterizes an ensemble, whereas f is (in this context) a stochastic variable, representing a property of a single gas:

$$(154) \quad f(\vec{r}, \vec{v}) = \frac{1}{N} \sum_i \delta(\vec{q}_i - \vec{r}) \delta(\vec{v} - \frac{\vec{p}_i}{m}).$$

How does $\rho_t^{(1)}$ evolve? From the Liouville-equation we get

$$(155) \quad \frac{\partial \rho^{(1)}(x_1)}{\partial t} = \int \{H, \rho\} d^3 \vec{p}_2 \cdots \vec{p}_N d\vec{q}_2 \cdots \vec{q}_N.$$

It is convenient here to regard the Poisson bracket as a differential operator on ρ , usually called the Liouville operator \mathcal{L} :

$$(156) \quad \mathcal{L}\rho := \sum_{i=1}^N \left(\frac{\partial H}{\partial \vec{q}_i} \cdot \frac{\partial}{\partial \vec{p}_i} - \frac{\partial H}{\partial \vec{p}_i} \cdot \frac{\partial}{\partial \vec{q}_i} \right) \rho.$$

For the Hamiltonian (149) this can be expanded as:

$$(157) \quad \mathcal{L} = \sum_i^N \mathcal{L}_i^{(1)} + \sum_{i < j}^N \mathcal{L}_{ij}^{(2)}$$

where

$$(158) \quad \mathcal{L}_i^{(1)} := \vec{p}_i \cdot \frac{\partial}{\partial \vec{q}_i}$$

and

$$(159) \quad \mathcal{L}_{ij}^{(2)} := \frac{\partial \phi_{ij}}{\partial \vec{q}_i} \cdot \left(\frac{\partial}{\partial \vec{p}_i} - \frac{\partial}{\partial \vec{p}_j} \right)$$

The evolution of $\rho^{(1)}$ is then given by:

$$(160) \quad \frac{\partial \rho_t^{(1)}(x_1)}{\partial t} = \mathcal{L}_1^{(1)} \rho_t^{(1)}(x_1) + \int dx_2 \mathcal{L}_{12}^{(2)} \rho^{(2)}(x_1, x_2)$$

More generally, for higher-order reduced distribution functions $\rho^{(m)}$ ($m \geq 2$), the evolution is governed by the equations:

$$(161) \quad \begin{aligned} \frac{\partial \rho_t^{(m)}(x_1, \dots, x_m)}{\partial t} &= \sum_{i=1}^m \mathcal{L}_i^{(1)} \rho_t^{(m)}(x_1, \dots, x_m) + \sum_{i < j=1}^m \mathcal{L}_{ij}^{(2)} \rho_t^{(m)}(x_1, \dots, x_m) \\ &+ \sum_i^m \int dx_{m+1} \mathcal{L}_{i, m+1}^{(2)} \rho_t^{(m+1)}(x_1, \dots, x_{m+1}) \end{aligned}$$

The equations (160,161) form the *BBGKY hierarchy*. It is strictly equivalent to the Hamiltonian formalism for symmetric ρ and H , provided that H contains no terms that depend on three or more particles. As one might expect, solving these equations is just as hard as for the original Hamiltonian equations. In particular, the equations are not closed: in order to know how $\rho_t^{(1)}$ evolves, we need to know $\rho_t^{(2)}$. In order to know how $\rho_t^{(2)}$ evolves, we need to know $\rho_t^{(3)}$ etc.

The usual method to overcome this problem is to cut off the hierarchy, i.e. to assume that for some finite m , $\rho^{(m)}$ is a functional of $\rho^{(\ell)}$ with $\ell < m$. In particular, if we just consider the easiest case ($m = 2$) and the easiest form of the functional, we can take $\rho^{(2)}$ to factorize in the distant past ($t \rightarrow -\infty$), giving:

$$(162) \quad \rho_t^{(2)}(x_1, x_2) = \rho_t^{(1)}(x_1) \rho_t^{(1)}(x_2); \quad \text{if } t \rightarrow -\infty$$

i.e., requiring that the molecular states of any pair of particles are uncorrelated *before* their interaction. This is analogous to the *Stoßzahlansatz* (29), but now, of course, formulated in terms of the reduced distribution functions of an ensemble.

It can be shown that for the homogeneous case, i.e. when $\rho^{(2)}$ is uniform over the positions \vec{q}_1 and \vec{q}_2 , i.e. when $\rho^{(2)}(x_1, x_2) = \rho^{(2)}(\vec{p}_1, \vec{p}_2)$ and when ϕ is a interaction potential of finite range, the evolution equation for $\rho^{(1)}$ becomes formally identical

to the Boltzmann equation (48). That is to say, in (160) we may substitute $\mathcal{L}_i^{(1)} = 0$ and:

$$\begin{aligned}
 \frac{\partial \rho_t^{(1)}(\vec{p}_1)}{\partial t} &= \int \mathcal{L}_{12}^{(2)} \rho(\vec{p}_1, \vec{p}_2) d^3 p_2 \\
 (163) \qquad &= \frac{N}{m} \int bdbd\phi \int d\vec{p}_2 \|\vec{p}_2 - \vec{p}_1\| \\
 &\qquad \left(\rho_t^{(1)}(\vec{p}_1') \rho_t^{(1)}(\vec{p}_2') - \rho_t^{(1)}(\vec{p}_1) \rho_t^{(1)}(\vec{p}_2) \right)
 \end{aligned}$$

(See [Uhlenbeck and Ford, 1963, p. 131] for more details.)

Remarks

The BBGKY approach is thoroughly Gibbsian in its outlook, i.e. it takes a probability density over phase space as its basic conceptual tool. An additional ingredient, not used extensively by Gibbs, is its reliance on permutation symmetry. It gives an enormous extension of Gibbs' own work by providing a systematic hierarchy of evolution equations for reduced (or marginalized) density functions, which can then be subjected to the techniques of perturbation theory. An ensemble-based analogy of the Boltzmann equation comes out of this approach as a first-order approximation for dilute gases with collision times much smaller than the mean free time. The Boltzmann equation for inhomogeneous gases cannot be obtained so easily— as one might expect also on physical grounds that one will need extra assumptions to motivate its validity.

It is instructive to compare this approach to Lanford's. His analogy of the Boltzmann equation is obtained for a different kind of function, namely the one-particle distribution function $F^{[x]}$, i.e. the exact relative number of particles with molecular state (\vec{r}, \vec{v}) , instead of $\rho^{(1)}$. Of course, there is a simple connection between the two. Noting that $F^{[x]}$ is a sum function (cf. equation (137)), we see that

$$(164) \quad \langle F^{[x]} \rangle = \int \rho^{(1)}(\vec{p}_1, \vec{q}_1) f(\delta(\vec{r} - \vec{q}_1) \delta(\vec{v} - \frac{\vec{p}_1}{m})) dp_1 dq_1 = \rho^{(1)}(\vec{r}, \vec{v}).$$

In other words, the one-particle distribution function $\rho^{(1)}$ is the expected value of the exact distribution of state. It thus appears that where Lanford describes the probability of the evolution of the exact distribution of state, the BBGKY result (164) describes the evolution of the average of the exact distribution of state. Lanford's results are therefore much more informative.

One might be tempted here to argue that one can justify or motivate that that actual particle distribution might be taken equal to its ensemble average by arguments similar to those employed in ergodic theory. In particular, we have seen from Khinchin's work (cf. §6.3) that for large enough systems, the probability that a sum function such as $F^{[x]}$ deviates significantly from its expectation value is

negligible. However, an important complication is that this reading of Khinchin's results holds for equilibrium, i.e. they apply with respect to the microcanonical distribution ρ_{mc} , not to an arbitrary time-dependent density ρ_t envisaged here.

The time asymmetry of the resulting equation does not derive from the hierarchy of equations, but from the ensemble-based analogy of the *Stoßzahlansatz* (162). That is to say, in this approach time asymmetry is introduced via an initial condition on the ensemble, i.e. the absence of initial correlations. It can be shown, just like for the original Boltzmann equation, that when the alternative boundary condition is imposed that makes the momenta independent after collisions, (i.e. if (162) is imposed for $t \rightarrow \infty$ instead) the anti-Boltzmann equation is obtained (see [Uhlenbeck and Ford, 1963, p. 127]).

7 STOCHASTIC DYNAMICS

7.1 Introduction

Over recent decades, some approaches to non-equilibrium statistical mechanics, that differ decidedly in their foundational and philosophical outlook, have nevertheless converged in developing a common unified mathematical framework. I will call this framework 'stochastic dynamics', since the main characteristic feature of the approach is that it characterizes the evolution of the state of a mechanical system as evolving under stochastic maps, rather than under a deterministic and time-reversal invariant Hamiltonian dynamics.⁶⁸

The motivations for adopting this stochastic type of dynamics come from different backgrounds, and one can find authors using at least three different views.

1. "Coarse graining" (cf. [van Kampen, 1962; Penrose, 1970]): In this view one assumes that on the microscopic level the system can be characterized as a (Hamiltonian) dynamical system with deterministic time-reversal invariant dynamics. However, on the macroscopic level, one is only interested in the evolution of macroscopic states, i.e. in a partition (or coarse graining) of the microscopic phase space into discrete cells. The usual idea is that the form and size of these cells are chosen in accordance with the limits of our observational capabilities. A more detailed exposition of this view is given in §7.5.

On the macroscopic level, the evolution now need no longer be portrayed as deterministic. When only the macrostate of a system at an instant is given, it is in general not fixed what its later macrostate will be, even if the underlying microscopic evolution is deterministic. Instead, one can provide *transition probabilities*, that specify how probable the transition from any given initial macrostate to later macrostates is. Although it is impossible, without further assumptions, to say anything general about the evolution of the macroscopically characterized states, it is possible to describe the evolution of an ensemble or a probability distribution over these states, in terms of a *stochastic process*.

⁶⁸Also, the name has been used in precisely this sense already by Sudarshan and coworkers, cf. [Sudarshan et al., 1961; Mehra and Sudarshan, 1972].

2. “Interventionism”, “tracing” or “open systems” (cf. [Blatt, 1959; Davies, 1976; Lindblad, 1976; Lindblad, 1983; Ridderbos, 2002]): On this view, one assumes that the system to be described is not isolated but in interaction with the environment. It is assumed that the total system, consisting of the system of interest and the environment can be described as a (Hamiltonian) dynamical system with a time-reversal invariant and deterministic dynamics. If we represent the state of the system by $x \in \Gamma^{(s)}$ and that of the environment by $y \in \Gamma^{(e)}$, their joint evolution is given by a one-parameter group of evolution transformations, generated from the Hamiltonian equations of motion for the combined system: $U_t : (x, y) \mapsto U_t(x, y) \in \Gamma^{(s)} \times \Gamma^{(e)}$. The evolution of the state x in the course of time is obtained by projecting, for each t , to the coordinates of $U_t(x, y)$ in $\Gamma^{(s)}$; call the result of this projection x_t . Clearly, this reduced time evolution of the system will generally fail to be deterministic, e.g. the trajectory described by x_t in $\Gamma^{(s)}$ may intersect itself.

Again, we may hope that this indeterministic evolution can nevertheless, for an ensemble of the system and its environment, be characterized as a stochastic process, at least if some further reasonable assumptions are made.

3. A third viewpoint is to deny [Mackey, 1992; 2001], or to remain agnostic about [Streater, 1995], the existence of an underlying deterministic or time-reversal invariant dynamics, and simply regard the evolution of a system as described by a stochastic process as a new fundamental form of dynamics in its own right.

While authors in this approach thus differ in their motivation and in the interpretation they have of its subject field, there is, as we shall see, a remarkable unity in the mathematical formalism adopted for this form of non-equilibrium statistical mechanics. The hope, obviously, is to arrange this description of the evolution of mechanical systems in terms of a stochastic dynamics in such a way that the evolution will typically display ‘irreversible behaviour’: i.e. an ‘approach to equilibrium’, that a Boltzmann-like evolution equation holds, that there is a stochastic analogy of the H -theorem, etc. In short, one would like to recover the autonomy and irreversibility that thermal systems in non-equilibrium states typically display.

We will see that much of this can be achieved with relatively little effort once a crucial technical assumption is in place: that the stochastic process is in fact a homogeneous Markov process, or, equivalently, obeys a so-called master equation. Much harder are the questions of whether the central assumptions of this approach might still be compatible with an underlying deterministic time-reversal invariant dynamics, and in which sense the results of the approach embody time-asymmetry. In fact we shall see that conflicting intuitions on this last issue arise, depending on whether one takes a probabilistic or a dynamics point of view towards this formalism.

From a foundational point of view, stochastic dynamics promises a new approach to the explanation of irreversible behaviour that differs in interesting ways from the more orthodox Hamiltonian or dynamical systems approach. In that approach, any account of irreversible phenomena can only proceed by referring to special initial conditions or dynamical hypotheses. Moreover, it is well-known that

an ensemble of such systems will conserve (fine-grained) Gibbs entropy so that the account cannot rely on this form of entropy for a derivation of the increase of entropy.

In stochastic dynamics, however, one may hope to find an account of irreversible behaviour that is not tied to special initial conditions, but one that is, so to say, built into the very stochastic-dynamical evolution. Further, since Liouville's theorem is not applicable, there is the prospect that one can obtain a genuine increase of Gibbs entropy from this type of dynamics.

As just mentioned, the central technical assumption in stochastic dynamics is that the processes described have the Markov property.⁶⁹ Indeed, general aspects of irreversible behaviour pour out almost effortlessly from the Markov property, or from the closely connected "master equation". Consequently, much of the attention in motivating stochastic dynamics has turned to the assumptions needed to obtain this Markov property, or slightly more strongly, to obtain a non-invertible Markov process [Mackey, 1992]. The best-known specimen of such an assumption is [van Kampen, 1962] "repeated randomness assumption". And similarly, critics of this type of approach [Sklar, 1993; Redhead, 1995; Callender, 1999] have also focused their objections on the question just how reasonable and general such assumptions are (cf. paragraph 7.5).

I believe both sides of the debate have badly missed the target. Many authors have uncritically assumed that the assumption of a (non-invertible) Markov process does indeed lead to non-time-reversal-invariant results. As a matter of fact, however, the Markov property (for invertible or non-invertible Markov processes) is time-reversal invariant. So, any argument to obtain that property need not presuppose time-asymmetry. In fact, I will argue that this discussion of irreversible behaviour as derived from the Markov property suffers from an illusion. It is due to the habit of studying the prediction of future states from a given initial state, rather than studying retrodictions towards an earlier state. As we shall see, for a proper description of irreversibility in stochastic dynamics one needs to focus on another issue, namely the difference between backward and forwards transition probabilities.

In the next subsections, I will first (§7.2) recall the standard definition of a homogeneous Markov process from the theory of stochastic processes. Subsection 7.3 casts these concepts in the language of dynamics, introduces the master equation, and discusses its analogy to the Boltzmann equation. In §7.4, we review some of the results that *prima facie* display irreversible behaviour for homogeneous Markov processes. In subsection 7.5 we turn to the physical motivations that have been given for the Markov property, and their problems, while §7.6 focuses on the question how seemingly irreversible results could have been obtained from a time-symmetric assumptions. Finally, §7.7 argues that a more promising discussion of these issues should start from a different definition of reversibility of stochastic processes.

⁶⁹Some authors argue that the approach can and should be extended to include non-Markovian stochastic processes as well. Nevertheless I will focus here on Markov processes.

7.2 The definition of Markov processes

To start off, consider an example. One of the oldest discussions of a stochastic process in the physics literature is the so-called ‘dog flea model’ of P. and T. Ehrenfest (1907).

Consider N fleas, labeled from 1 to N , situated on either of two dogs. The number of fleas on dog 1 and 2 are denoted as n_1 and $n_2 = N - n_1$. Further, we suppose there is an urn with N lots carrying the numbers $1, \dots, N$ respectively. The urn is shaken, a lot is drawn (and replaced), and the flea with the corresponding label is ordered to jump to the other dog. This procedure is repeated every second.

It is not hard to see that this model embodies an ‘approach equilibrium’ in some sense: Suppose that initially all or almost all fleas are on dog 1. Then it is very probable that the first few drawings will move fleas from dog 1 to 2. But as soon as the number of fleas on dog 2 increases, the probability that some fleas will jump back to dog 1 increases too. The typical behaviour of, say, $|n_1 - n_2|$ as a function of time will be similar to Boltzmann’s H -curve, with a tendency of $|n_1 - n_2|$ to decrease if it was initially large, and to remain close to the ‘equilibrium’ value $n_1 \approx n_2$ for most of the time. But note that in contrast to Boltzmann’s H -curve in gas theory, the ‘evolution’ is here entirely stochastic, i.e. generated by a lottery, and that no underlying deterministic equations of motion are provided.

In general, a stochastic process is, mathematically speaking, nothing but a probability measure P on a measure space X , whose elements will here be denoted as ξ , on which there are infinitely many random variables Y_t , with $t \in \mathbb{R}$ (or sometimes $t \in \mathbb{Z}$). Physically speaking, we interpret t as time, and Y as the macroscopic variable(s) characterizing the macrostate — say the number of fleas on a dog, or the number of molecules with their molecular state in some cell of μ -space, etc. Further, ξ represents the total history of the system which determines the values of $Y_t(\xi)$. The collection Y_t may thus be considered as a single random variable Y evolving in the course of time.

At first sight, the name ‘process’ for a probability measure may seem somewhat unnatural. From a physical point of view it is the *realization*, in which the random variables Y_t attain the values $Y_t(\xi) = y_t$ that should be called a process. In the mathematical literature, however, it has become usual to denote the measure that determines the probability of all such realizations as a ‘stochastic process’.

For convenience we assume here that the variables Y_t may attain only finitely many discrete values, say $y_t \in \mathcal{Y} = \{1, \dots, m\}$. However, the theory can largely be set up in complete analogy for continuous variables.

The probability measure P provides, for $n = 1, 2, \dots$, and instants t_1, \dots, t_n definite probabilities for the event that Y_t at these instants attains certain values y_1, \dots, y_n :

$$\begin{aligned} P_{(1)}(y_1, t_1) \\ P_{(2)}(y_2, t_2; y_1, t_1) \\ \vdots \end{aligned}$$

$$(165) \quad \begin{aligned} &P_{(n)}(y_n, t_n; \dots; y_1, t_1) \\ &\quad \vdots \end{aligned}$$

Here, $P_{(n)}(y_n, t_n; \dots; y_1, t_1)$ stands for the joint probability that at times t_1, \dots, t_n the quantities Y_t attain the values y_1, \dots, y_n , with $y_i \in \mathcal{Y}$. It is an abbreviation for

$$(166) \quad P_{(n)}(y_n, t_n; \dots; y_1, t_1) := P(\{\xi \in X : Y_{t_n}(\xi) = y_n \ \& \ \dots \ \& \ Y_{t_1}(\xi) = y_1\})$$

Obviously the probabilities (165) are normalized and non-negative, and each $P_{(n)}$ is a marginal of all higher-order probability distributions:

$$(167) \quad P_{(n)}(y_n, t_n; \dots; y_1, t_1) = \sum_{y_{n+m}} \dots \sum_{y_{n+1}} P_{(n+m)}(y_{n+m}, t_{n+m}; \dots; y_1, t_1).$$

In fact, the probability measure P is uniquely determined by the hierarchy (165).⁷⁰

Similarly, we may define conditional probabilities in the familiar manner, e.g.:

$$(168) \quad P_{(1|n-1)}(y_n, t_n | y_{n-1}, t_{n-1}; \dots; y_1, t_1) := \frac{P_{(n)}(y_n, t_n; \dots; y_1, t_1)}{P_{(n-1)}(y_{n-1}, t_{n-1}; \dots; y_1, t_1)}$$

provides the probability that Y_{t_n} attains the value y_n , under the condition that $Y_{t_{n-1}}, \dots, Y_{t_1}$ have the values y_{n-1}, \dots, y_1 .

In principle, the times appearing in the joint and conditional probability distributions (165,168) may be chosen in an arbitrary order. However, we adopt from now on the convention that they are ordered as $t_1 < \dots < t_n$.

A special and important type of stochastic process is obtained by adding the assumption that such conditional probabilities depend only the condition at the last instant. That is to say: for all n and all choices of y_1, \dots, y_n and $t_1 < \dots < t_n$:

$$(169) \quad P_{(1|n)}(y_n, t_n | y_{n-1}, t_{n-1}; \dots; y_n, t_n) = P_{(1|1)}(y_n, t_n | y_{n-1}, t_{n-1})$$

This is the *Markov property* and such stochastic processes are called *Markov processes*.

The interpretation often given to this assumption, is that Markov processes have ‘no memory’. To explain this slogan more precisely, consider the following situation. Suppose we are given a piece of the history of the quantity Y : at the instants t_1, \dots, t_{n-1} its values have been y_1, \dots, y_{n-1} . On this information, we want to make a prediction of the value y_n of the variable Y at a later instant t_n . The Markov-property (169) says that this prediction would not have been better or worse if, instead of knowing this entire piece of prehistory, only the value y_{n-1} of Y at the last instant t_{n-1} had been given. Additional information about the past values is thus irrelevant for a prediction of the future value.

For a Markov process, the hierarchy of joint probability distributions (165) is subjected to stringent demands. In fact they are all completely determined by:

⁷⁰At least, when we assume that the σ -algebra of measurable sets in X is the cylinder algebra generated by sets of the form in the right-hand side of (166).

- (a) the specification of $P_{(1)}(y, 0)$ at one arbitrary chosen initial instant $t = 0$, and
- (b) the conditional probabilities $P_{(1|1)}(y_2, t_2|y_1, t_1)$ for all $t_2 > t_1$. Indeed,

$$(170) \quad P_{(1)}(y, t) = \sum_{y_0} P_{(1|1)}(y, t|y_0, 0)P_{(1)}(y_0, 0);$$

and for the joint probability distributions $P_{(n)}$ we find:

$$(171) \quad P_{(n)}(y_n, t_n; \dots; y_1, t_1) = P_{(1|1)}(y_n, t_n|y_{n-1}, t_{n-1})P_{(1|1)}(y_{n-1}, t_{n-1}|y_{n-2}, t_{n-2}) \times \\ \times \dots \times P_{(1|1)}(y_2, t_2|y_1, t_1)P_{(1)}(y_1, t_1).$$

It follows from the Markov property that the conditional probabilities $P_{(1|1)}$ have the following property, known as the *Chapman-Kolmogorov* equation:

$$(172) \quad P_{(1|1)}(y_3, t_3|y_1, t_1) = \sum_{y_2} P_{(1|1)}(y_3, t_3|y_2, t_2)P_{(1|1)}(y_2, t_2|y_1, t_1) \quad \text{for } t_1 < t_2 < t_3.$$

So, for a Markov process, the hierarchy (165) is completely characterized by specifying $P_{(1)}$ at an initial instant and a system of conditional probabilities $P_{(1|1)}$ satisfying the Chapman-Kolmogorov equation. The study of Markov processes therefore focuses on these two ingredients.⁷¹

A following special assumption is *homogeneity*. A Markov process is called homogeneous if the conditional probabilities $P_{(1|1)}(y_2, t_2|y_1, t_1)$ do not depend on the two times t_1, t_2 separately but only on their mutual difference $t = t_2 - t_1$; i.e. if they are invariant under time translations. In this case we may write

$$(173) \quad P_{(1|1)}(y_2, t_2|y_1, t_1) = T_t(y_2, y_1)$$

such conditional probabilities are also called *transition* probabilities.

Is the definition of a Markov process time-symmetric? The choice in (169) of conditionalizing the probability distribution for Y_{t_n} on *earlier* values of Y_t is of course special. In principle, there is nothing in the formulas (165) or (168) that forces such an ordering. One might, just as well, ask for the probability of a value of Y_t in the past, under the condition that part of the *later* behaviour is given (or, indeed, conditionalize on the behaviour at both earlier and later instants.)

At first sight, the Markov property makes no demands about these latter cases. Therefore, one might easily get the impression that the definition is time-asymmetric. However, this is not the case. One can show that (169) is equivalent to:

$$(174) \quad P_{(1|n-1)}(y_1, t_1|y_2, t_2; \dots; y_n, t_n) = P_{(1|1)}(y_1 t_1|y_2, t_2)$$

⁷¹Note, however, that although every Markov process is fully characterized by (i) an initial distribution $P_{(1)}(y, 0)$ and (ii) a set of transition probabilities $P_{(1|1)}$ obeying the Chapman-Kolmogorov equation and the equations (171), it is *not* the case that every stochastic process obeying (i) and (ii) is a Markov process. (See [van Kampen, 1981, p. 83] for a counterexample). Still, it is true that one can define a unique Markov process from these two ingredients by stipulating (171).

where the convention $t_1 < t_2 < \dots < t_n$ is still in force. Thus, a Markov process does not only have ‘no memory’ but also ‘no foresight’. Some authors (e.g. [Kelly, 1979]) adopt an (equivalent) definition of a Markov process that is explicitly time-symmetric: Suppose that the value y_i at an instant t_i somewhere in the middle of the sequence $t_1 < \dots < t_n$ is given. The condition for a stochastic process to be Markov is then

$$(175) \quad P_{(n|1)}(y_n, t_n; \dots; y_1, t_1 | y_i, t_i) = P_{(n-i|1)}(y_n, t_n; \dots; y_{i+1}, t_{i+1} | y_i, t_i) P_{(i-1|1)}(y_{i-1}, t_{i-1}; y_1, t_1 | y_i, t_i)$$

for all $n = 1, 2, \dots$ and all $1 \leq i \leq n$. In another slogan: The future and past are independent if one conditionalizes on the present.

7.3 Stochastic dynamics

A homogeneous Markov process is for $t > 0$ completely determined by the specification of an initial probability distribution $P_{(1)}(y, 0)$ and the transition probabilities $T_t(y_2|y_1)$ defined by (173). The difference in notation (between P and T) also serves to ease a certain conceptual step. Namely, the idea is to regard T_t as a stochastic evolution operator. Thus, we can regard $T_t(y_2|y_1)$ as the elements of a matrix, representing a (linear) operator T that determines how an initial distribution $P_{(1)}(y, 0)$ will evolve into a distribution at later instants $t > 0$. (In the sequel I will adapt the notation and write $P_{(1)}(y, t)$ as $P_t(y)$.)

$$(176) \quad P_t(y) = (T_t P)(y) := \sum_{y'} T_t(y|y') P_0(y').$$

The Chapman-Kolmogorov equation (172) may then be written compactly as

$$(177) \quad T_{t+t'} = T_t \circ T_{t'} \quad \text{for } t, t' \geq 0$$

where \circ stands for matrix multiplication, and we now also extend the notation to include the unit operator:

$$(178) \quad \mathbf{1}(y, y') = T_0(y, y') := \delta_{y, y'}$$

where δ denotes the Kronecker delta.

The formulation (177) can (almost) be interpreted as the group composition property of the evolution operators T . It may be instructive to note how much this is due to the Markov property. Indeed, for arbitrary conditional probabilities, say, if A_i, B_j and C_k denote three families of complete and mutually exclusive events (i.e. $\cup_i A_i = \cup_j B_j = \cup_k C_k = \mathcal{Y}$; $A_i \cap A_{i'} = B_j \cap B_{j'} = C_k \cap C_{k'} = \emptyset$ for $i \neq i', j \neq j'$ and $k \neq k'$), the rule of total probability gives :

$$(179) \quad P(A_i|C_k) = \sum_j P(A_i|B_j, C_k) P(B_j|C_k).$$

In general, this rule can *not* be regarded as ordinary matrix multiplication or a group composition! But the Markov property makes $P(A_i|B_j, C_k)$ in (179) reduce to $P(A_i|B_j)$, and then the summation in (179) coincides with familiar rule for matrix multiplication.

I wrote above: ‘almost’, because there is still a difference in comparison with the normal group property: in the Chapman-Kolmogorov-equation (177) all times must be positive. Thus, in general, for $t > 0$, T_t may not even be defined and so it does *not* hold that

$$(180) \quad T_{-t} \circ T_t = \mathbf{1}.$$

A family of operators $\{T_t, t \geq 0\}$ which is closed under a operation \circ that obeys (177), and for which $T_0 = \mathbf{1}$ is called a *semigroup*. It differs from a group in the sense that its elements T_t need not be *invertible*, i.e., need not have an inverse. The lack of an inverse of T_t may be due to various reasons: either T_t does not possess an inverse, i.e. it is not a one-to-one mapping, or T_t does possess an inverse matrix T_t^{inv} , which however is itself non-stochastic (e.g. it may have negative matrix-elements). We will come back to the role of the inverse matrices in Sections 7.4 and 7.7.

The theory of Markov processes has a strong and natural connection with linear algebra. Sometimes, the theory is presented entirely from this perspective, and one starts with the introduction of a semigroup of *stochastic matrices*, that is to say, m by m matrices T with $T_{ij} \geq 0$ and $\sum_i T_{ij} = 1$. Or, more abstractly, one posits a class of states P , elements of a Banach space with a norm $\|P\|_1 = 1$, and a semigroup of stochastic maps T_t , ($t \geq 0$), subject to the conditions that T_t is linear, positive, and preserves norm: $\|T_t P\|_1 = \|P\|_1$, (cf. [Streater, 1995]).

The evolution of a probability distribution P (now regarded as a vector or a state) is then particularly simple when t is discrete ($t \in \mathbb{N}$):

$$(181) \quad P_t = T^t P_0, \quad \text{where } T^t = \underbrace{T \circ \dots \circ T}_{t \text{ times}}.$$

Homogeneous Markov processes in discrete time are also known as *Markov chains*.

Clearly, if we consider the family $\{T_t\}$ as a semigroup of stochastic evolution operators, or a stochastic form of dynamics, it becomes attractive to look upon $P_0(y)$ as a contingent initial state, chosen independently of the evolution operators T_t . Still, from the perspective of the probabilistic formalism with which we started, this might be an unexpected thought: both $P_{(1)}$ and $P_{(1|1)}$ are aspects of a single, given, probability measure P . The idea of regarding them as independent ingredients that may be specified separately doesn’t then seem very natural. But, of course, there is no formal objection against the idea, since every combination of a system of transition probabilities T_t obeying the Chapman-Kolmogorov equation, and an arbitrary initial probability distribution $P_0(y) = P_{(1)}(y, 0)$ defines a unique homogeneous Markov process (cf. footnote 71). In fact, one sometimes even goes one step further and identifies a homogeneous Markov process completely with the specification of the transition probabilities, without regard of the initial state

$P_0(y)$; just like the dynamics of a deterministic system is usually presented without assuming any special initial state.

For Markov chains, the goal of specifying the evolution of $P_t(y)$ is now already completely solved in equation (181). In the case of continuous time, it is more usual to specify evolution by means of a differential equation. Such an equation may be obtained in a straightforward manner by considering a Taylor expansion of the transition probability for small times [van Kampen, 1981, p.101–103] — under an appropriate continuity assumption.

The result (with a slightly changed notation) is:

$$(182) \quad \frac{\partial P_t(y)}{\partial t} = \sum_{y'} (W(y|y')P_t(y') - W(y'|y)P_t(y))$$

Here, the expression $W(y|y')$ is the transition probability from y' to y *per unit of time*. This differential equation, first obtained by Pauli in 1928, is called the *master equation*. (This name has become popular because an equation of this type covers a great variety of processes.)

The interpretation of the equation is suggestive: the change of the probability $P_t(y)$ is determined by making up a balance between gains and losses: the probability of value y increases in a time dt because of the transitions from y' to y , for all possible values of y' . This increase per unit of time is $\sum_{y'} W(y|y')P_t(y')$. But in same period dt there is also a decrease of $P_t(y)$ as a consequence of transitions from the value y to all other possible values y' . This provides the second term.

In this “balancing” aspect, the master equation resembles the Boltzmann equation (48), despite the totally different derivation, and the fact that $P_t(y)$ has quite another meaning than Boltzmann’s $f_t(v)$. (The former is a probability distribution, the latter a distribution of particles.) Both are first-order differential equations in t . A crucial mathematical distinction from the Boltzmann equation is that the master equation is linear in P , and therefore much easier to solve.

Indeed, any solution of the master equation can formally be written as:

$$(183) \quad P_t = e^{tL} P_0,$$

where L represents the operator

$$(184) \quad L(y|y') := W(y|y') - \sum_{y''} W(y''|y')\delta_{y,y''}.$$

The general solution (183) is similar to the discrete time case (181), thus showing the equivalence of the master equation to the assumption of a homogeneous Markov process in continuous time.

A final remark(not needed for later paragraphs). The analogy with the Boltzmann equation can even be increased by considering a Markov process for particle pairs, i.e. by imagining a process where pairs of particles with initial states (i, j) make a transition to states (k, l) with certain transition probabilities (cf. [Alberti and Uhlmann, 1982, p. 30]) Let $W(i, j|k, l)$ denote the associated transition probability per unit of time. Then the master equation takes the form:

$$(185) \quad \frac{\partial P_t(i, j)}{\partial t} = \sum_{k, l} (W(i, j|k, l)P_t(k, l) - W(k, l|i, j)P_t(i, j)).$$

Assume now that the transitions $(i, j) \rightarrow (k, l)$ and $(k, l) \rightarrow (i, j)$ are equally probable, so that the transition probability per unit of time is symmetric: $W(i, j|k, l) = W(k, l|i, j)$, and, as an analogue to the *Stoßzahlansatz*, that $P(i, j)$ in the right-hand side may be replaced by the product of its marginals:

$$(186) \quad P(i, j) \rightarrow \sum_j P(i, j) \cdot \sum_i P(i, j) = P'(i)P''(j)$$

Summing the above equation (185) over j , we finally obtain

$$(187) \quad \frac{\partial P'_t(i)}{\partial t} = \sum_j \frac{\partial P_t(i, j)}{\partial t} = \sum_{j, k, l} T(i, j|k, l)(P'_t(k)P''_t(l) - P'_t(i)P''_t(j)),$$

i.e., an even more striking analogue of the Boltzmann equation (48). But note that although (185) describes a Markov process, the last equation (187) does not: it is no longer linear in P , as a consequence of the substitution (186).

7.4 Approach to equilibrium and increase of entropy?

What can we say in general about the evolution of $P_t(y)$ for a homogeneous Markov process? An immediate result is this: the *relative entropy* is monotonically non-decreasing. That is to say, if we define

$$(188) \quad H(P, Q) := - \sum_{y \in \mathcal{Y}} P(y) \ln \frac{P(y)}{Q(y)}$$

as the relative entropy of a probability distribution P relative to Q , then one can show (see e.g. Moran 1961; Mackey 1991, p. 30):

$$(189) \quad H(P_t, Q_t) \geq H(P, Q)$$

where $P_t = T_t P$, $Q_t = T_t Q$, and T_t are elements of the semigroup (181) or (183).

One can also show that a non-zero relative entropy increase for at least some pair probability distributions P and Q , the stochastic matrix T_t must be non-invertible.

The relative entropy $H(P|Q)$ can, in some sense, be thought of as a measure of how much P and Q “resemble” each other.⁷² Indeed, it takes its maximum value (i.e. 0) if and only if $P = Q$; it may become $-\infty$ if P and Q have disjoint support, (i.e. when $P(y)Q(y) = 0$ for all $y \in \mathcal{Y}$.) Thus, the result (189) says that if the stochastic process is non-invertible, pairs of distributions P_t and Q_t will generally become more and more alike as time goes by.

Hence it seems we have obtained a general weak aspect of “irreversible behaviour” in this framework. Of course, the above result does not yet imply that

⁷²Of course, this is an asymmetric sense of “resemblance” because $H(P, Q) \neq H(Q, P)$.

the ‘absolute’ entropy $H(P) := -\sum_y P(y) \ln P(y)$ of a probability distribution is non-decreasing. But now assume that the process has a *stationary state*. In other words, there is a probability distribution $P^*(y)$ such that

$$(190) \quad T_t P^* = P^*.$$

The intention is, obviously, to regard such a distribution as a candidate for the description of an equilibrium state. If there is such a stationary distribution P^* , we may apply the previous result and write:

$$(191) \quad H(P, P^*) \leq H(T_t P, T_t P^*) = H(P_t, P^*).$$

In other words, as time goes by, the distribution $T_t P$ will then more and more resemble the stationary distribution than does P . If the stationary distribution is also uniform, i.e.:

$$(192) \quad P^*(y) = \frac{1}{m},$$

then not only the relative but also the absolute entropy $H(P) := -\sum_y P(y) \ln P(y)$ increases, because

$$(193) \quad H(P, P^*) = H(P) - \ln m.$$

In order to get a satisfactory description of an ‘approach to equilibrium’ the following questions remain:

- (i) is there such a stationary distribution?
- (ii) if so, is it unique?
- (iii) does the monotonic behaviour of $H(P_t)$ imply that $\lim_{t \rightarrow \infty} P_t = P^*$?

Harder questions, which we postpone to the next subsection 7.5, are:

- (iv) how to motivate the assumptions needed in this approach or how to make judge their (in)compatibility with an underlying time deterministic dynamics; and
- (v) how this behaviour is compatible with the time symmetry of Markov processes.

Ad (i). A stationary state as defined by (190), can be seen as an eigenvector of T_t with eigenvalue 1, or, in the light of (183), an eigenvector of L for the eigenvalue 0. Note that T or L are not necessarily Hermitian (or, rather, since we are dealing with real matrices, symmetric), so that the existence of eigenvectors is not guaranteed by the spectral theorem. Further, even if an eigenvector with the corresponding eigenvalue exists, it is not automatically suitable as a probability distribution because its components might not be positive.

Still, it turns out that, due to a theorem of Perron (1907) and Frobenius (1909), every stochastic matrix indeed has a eigenvector, with exclusively non-negative components, and eigenvalue 1 (see e.g. [Gantmacher, 1959; Van Harn and Holewijn, 1991]). But if the set \mathcal{Y} is infinite or continuous this is not always true.

A well-known example of the latter case is the so-called Wiener process that is often used for the description of Brownian motion. It is characterized by the transition probability density:

$$(194) \quad T_t(y|y') = \frac{1}{\sqrt{2\pi t}} \exp \frac{(y - y')^2}{2t}, \quad y, y' \in \mathbb{R}.$$

The evolution of an arbitrary initial probability density ρ_0 can be written as a convolution:

$$(195) \quad \rho_t(y) = \int T_t(y|y')\rho_0(y')dy';$$

which becomes gradually lower, smoother and wider in the course of time, but does not approach any stationary probability density. Because this holds for every choice of ρ_0 , there is no stationary distribution in this case.

However, it is not reasonable to see this as a serious defect. Indeed, in thermodynamics too one finds that a plume of gas emitted into free space will similarly diffuse, becoming ever more dilute without ever approaching an equilibrium state. Thermodynamic equilibrium is only approached for systems enclosed in a vessel of finite volume.

However, for continuous variables with a range that has finite measure, the existence of a stationary distribution is guaranteed under the condition that the probability density ρ_y is at all times bounded, i.e. $\exists M \in \mathbb{R}$ such that $\forall t \rho_t \leq M$; (see [Mackey, 1992, p. 36]).

Ad (ii). The question whether stationary solutions will be unique is somewhat harder to tackle. This problem exhibits an analogy to that of metric transitivity in the ergodic problem (cf. paragraph 6.1).

In general, it is very well possible that the range \mathcal{Y} of Y can be partitioned in two disjoint regions, say A and B , with $\mathcal{Y} = A \cup B$, such that there are no transitions from A to B or vice versa (or that such transitions occur with probability zero). That is to say, the stochastic evolution T_t might have the property

$$(196) \quad T_t(Y \in A|Y \in B) = T_t(Y \in B|Y \in A) = 0$$

In other words, its matrix may, (perhaps after a conventional relabeling of the outcomes) be written in the form:

$$(197) \quad \begin{pmatrix} T_A & 0 \\ 0 & T_B \end{pmatrix}.$$

The matrix is then called (completely) *reducible*. In this case, stationary distributions will generally not be unique: If P_A^* is a stationary distribution with support in the region A , and P_B^* is a stationary distribution with support in B , then every convex combination

$$(198) \quad \alpha P_A^*(y) + (1 - \alpha)P_B^*(y) \quad \text{with } 0 \leq \alpha \leq 1.$$

will be stationary too. In order to obtain a unique stationary solution we will thus have to assume an analogue of metric transitivity. That is to say: we should

demand that every partition of \mathcal{Y} into disjoint sets A and B for which (196) holds is ‘trivial’ in the sense that $P(A) = 0$ or $P(B) = 0$.

So, one may ask, is the stationary distribution P^* unique if and only if the transition probabilities T_τ are not reducible? In the ergodic problem, as we saw in 6.1, the answer is positive (at least if P^* is assumed to be absolutely continuous with respect to the microcanonical measure). But not in the present case!

This has to do with the phenomenon of so-called ‘transient states’, which has no analogy in Hamiltonian dynamics. Let us look at an example to introduce this concept. Consider a stochastic matrix of the form:

$$(199) \quad \begin{pmatrix} T_A & T' \\ 0 & T_B \end{pmatrix}$$

where T' is a matrix with non-negative entries only. Then:

$$(200) \quad \begin{pmatrix} T_A & T' \\ 0 & T_B \end{pmatrix} \begin{pmatrix} P_A \\ 0 \end{pmatrix} = \begin{pmatrix} T_A P_A \\ 0 \end{pmatrix}, \quad \begin{pmatrix} T_A & T' \\ 0 & T_B \end{pmatrix} \begin{pmatrix} 0 \\ P_B \end{pmatrix} = \begin{pmatrix} T' P_B \\ T_B P_B \end{pmatrix}$$

so that here transitions of the type $a \rightarrow b$ have probability zero, but transitions of the type $b \rightarrow a$ occur with positive probability. (Here, a, b stand for arbitrary elements of the subsets A and B .) It is clear that in such a case the region B will eventually be ‘sucked empty’. That is to say: the total probability of being in region B (i.e. $\|T^t P_B\|$) will go exponentially to zero. The distributions with support in B are called ‘transient’ and the set A is called ‘absorbing’ or a ‘trap’.

It is clear that these transient states will not play any role in the determination of the stationary distribution, and that for this purpose they might be simply ignored. Thus, in this example, the only stationary states are those with a support in A . And there will be more than one of them if T_A is reducible.

A matrix T that may be brought (by permutation of the rows and columns) in the form (199), with T_A reducible is called *incompletely reducible* [van Kampen, 1981, p. 108]. Further, a stochastic matrix is called *irreducible* if it is neither completely or incompletely reducible. An alternative (equivalent) criterion is that all states ‘communicate’ with each other, i.e. that for every pair of $i, j \in \mathcal{Y}$ there is some time t such that $P_t(j|i) > 0$.

The Perron-Frobenius theorem guarantees that as long as T irreducible, there is a unique stationary distribution. Furthermore, one can then prove an analogue of the ergodic theorem:[Petersen, 1983, p. 52]

ERGODIC THEOREM FOR MARKOV PROCESSES: If the transition probability T_t is irreducible, the time average of P_t converges to the unique stationary solution:

$$(201) \quad \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau T_t P(y) dt = P^*(y).$$

Ad (iii). If there is a unique stationary distribution P^* , will $T_t P$ converge to P^* , for every choice of P ? Again, the answer is not necessarily affirmative. (Even if (201) is valid!) For example, there are homogeneous and irreducible Markov chains for which P_t can be divided into two pieces: $P_t = Q_t + R_t$ with the following properties [Mackey, 1992, p. 71]:

1. Q_t is a term with $\|Q_t\| \rightarrow 0$. This is a transient term.
2. The remainder R_t is periodic, i.e. after some finite time τ the evolution repeats itself: $R_{t+\tau} = R_t$.

These processes are called *asymptotically periodic*. They may very well occur in conjunction with a unique stationary distribution P^* , and show strict monotonic increase of entropy, but still not converge to P^* . In this case, the monotonic increase of relative entropy $H(P_t, P^*)$ is entirely due to the transient term. For the periodic piece R_t , the transition probabilities are permutation matrices, which, after τ repetitions, return to the unit operator.

Besides, if we arrange that P^* is uniform, we can say even more in this example: The various forms R_t that are attained during the cycle of permutations with period τ all have the same value for the relative entropy $H(R_t, P^*)$, but this entropy is strictly lower than $H(P^*, P^*) = 0$. In fact, P^* is the average of the R_t 's, i.e.: $P^* = \frac{1}{\tau} \sum_{t=1}^{t=\tau} R_t$, in correspondence with (201).

Further technical assumptions can be introduced to block examples of this kind, and thus enforce a strict convergence towards the unique stationary distribution, e.g. by imposing a condition of 'exactness' [Mackey, 1992]. However, it would take us too far afield to discuss this in detail.

In conclusion, it seems that a weak aspect of "irreversible behaviour", i.e. the monotonic non-decrease of relative entropy is a general feature for all homogeneous Markov processes, (and indeed for all stochastic processes), and non-trivially so when the transition probabilities are non-invertible. Stronger versions of that behaviour, in the sense of affirmative answers to the questions (i), (ii) and (iii), can be obtained too, but at the price of additional technical assumptions.

7.5 Motivations for the Markov property and objections against them

Ad (iv). We now turn to the following problem: what is the motivation behind the assumption of the Markov property? The answer, of course, is going to depend on the interpretation of the formalism that one has in mind, and may be different in the 'coarse-graining' and the 'open systems' or interventionist approaches (cf. Section 7.1). I shall discuss the coarse-graining approach in the next paragraph below, and then consider the similar problem for the interventionist point of view.

Coarse-graining and the repeated randomness assumption

In the present point of view, one assumes that the system considered is really an isolated Hamiltonian system, but the Markov property is supposedly obtained from a partitioning of its phase space. But exactly how is that achieved?

One of the clearest and most outspoken presentations of this view is [van Kampen, 1962]. As in paragraph 5.4, we assume the existence of some privileged partition of the Hamiltonian phase space Γ — or of the energy hypersurface Γ_E — into disjoint cells: $\Gamma = \omega_1 \cup \dots \cup \omega_m$. Consider an arbitrary ensemble with probability density ρ on this phase space. Its evolution can be represented by an operator

$$(202) \quad U_t^* \rho(x) := \rho(U_{-t}x),$$

where, — in order to avoid conflation of notation — we now use U_t to denote the Hamiltonian evolution operators, previously denoted as T_t , e.g. in (68) and throughout section 6. Let transition probabilities between the cells of this partition be defined as

$$(203) \quad T_t(j|i) := P(x_t \in \omega_j | x \in \omega_i) = P(U_t x \in \omega_j | x \in \omega_i) = \frac{\int_{(U_{-t}\omega_j) \cap \omega_i} \rho(x) dx}{\int_{\omega_i} \rho(x) dx},$$

Obviously such transition probabilities will be homogeneous, due to the time-translation invariance of the Hamiltonian evolution U_t . Further, let $\hat{p}_0(i) := P(x \in \omega_i) = \int_{\omega_i} \rho(x) dx$, $i \in \mathcal{Y} = \{1, \dots, m\}$, be an arbitrary initial coarse-grained probability distribution at time $t=0$.

Using the coarse-graining map defined by (92), one may also express the coarse-grained distribution at time t as

$$(204) \quad \mathcal{C}G U_t^* \rho(x) = \sum_{ji} T_t(j|i) \hat{p}_0(i) \frac{1}{\mu(\omega_j)} \mathbf{1}_{\omega_j}(x)$$

where μ is the canonical measure on Γ , or the microcanonical measure on Γ_E . This expression indicates that, as long as we are only interested in the coarse grained history, it suffices to know the transition probabilities (203) and the initial coarse grained distributions.

But in order to taste the fruits advertised in the previous paragraphs, one needs to show that the transition probabilities define a Markov process, i.e., that they obey the Chapman-Kolmogorov equation (172),

$$(205) \quad T_{t'+t}(k|i) = T_{t'}(k|j)T_t(j|i); \quad \text{for all } t, t' > 0.$$

Applying (204) for times t , t' and $t + t'$, it follows easily that the Chapman-Kolmogorov equation is equivalent to

$$(206) \quad \mathcal{C}G U_{t'+t}^* = \mathcal{C}G U_{t'}^* \mathcal{C}G U_t^*, \quad \text{for all } t, t' > 0.$$

In other words, the coarse-grained probability distribution at time $t + t'$ can be obtained by first applying the Hamiltonian dynamical evolution during a time t , then performing a coarse-graining operation, next applying the dynamical evolution during time t' , and then coarse-graining again. In comparison to the relation $U_{t'+t}^* = U_t^* U_t^*$, we see that the Chapman-Kolmogorov condition can be obtained by demanding that it is allowed to apply a coarse-graining, i.e. to reshuffle the phase points within each cell at any intermediate stage of the evolution. Of course, this coarse-graining halfway during the evolution erases all information about the past evolution apart from the label of the cell where the state is located at that time; and this ties in nicely with the view of the Markov property as having no memory (cf. the discussion on p. 1042).

What is more, the *repeated* application of the coarse-graining does lead to a monotonic non-decrease of the Gibbs entropy: If, for simplicity, we divide a time interval into m segments of duration τ , we have

$$(207) \quad \rho_{m\tau} = \underbrace{\mathcal{C}G U_\tau^* \mathcal{C}G U_\tau^* \cdots \mathcal{C}G U_\tau^*}_{m \text{ times}} \rho$$

and from (96):

$$(208) \quad \sigma[\rho_{m\tau}] \geq \sigma[\rho_{(m-1)\tau}] \geq \cdots \geq \sigma[\rho_\tau] \geq \sigma[\rho_0].$$

But since the choice of τ is arbitrary, we may conclude that $\sigma[\rho_t]$ is monotonically non-decreasing.

Thus, van Kampen argues, the ingredient to be added to the dynamical evolution is that, at any stage of the evolution, one should apply a coarse-graining of the distribution. It is important to note that it is not sufficient to do that just once at a single instant. At every stage of the evolution we need to coarse-grain the distribution again and again. Van Kampen [1962, p. 193] calls this the *repeated randomness* assumption.

What is the justification for this assumption? Van Kampen points out that it is “not unreasonable” (ibid., p. 182), because of the brute fact of its success in phenomenological physics. Thermodynamics and other phenomenological descriptions of macroscopic systems (the diffusion equation, transport equations, hydrodynamics, the Fokker-Planck equation, etc.) all characterize macroscopic systems with a very small number of variables. This means that their state descriptions are very coarse in comparison with the microscopic phase space. But their evolution equations are autonomous and deterministic: the change of the macroscopic variables is given in terms of the instantaneous values of those very same variables. The success of these equations shows, apparently, that the precise microscopic state does not add any relevant information beyond this coarse description. At the same time, van Kampen admits that the coarse-graining procedure is clearly not always successful. It is not difficult to construct a partition of a phase space into cells for which the Markov property fails completely.

Apparently, the choice of the cells must be “just right” [van Kampen, 1962, p. 183]. But there is as yet no clear prescription how this is to be done. Van

Kampen [1981, p. 80] argues that it is “the art of the physicist” to find the right choice, an art in which he or she succeeds in practice by a mixture of general principles and ingenuity, but where no general guidelines can be provided. The justification of the repeated randomness assumption is that it leads to the Markov property and from there onwards to the master equation, providing a successful autonomous, deterministic description of the evolution of the coarse-grained distribution.

It is worth noting that van Kampen thus qualifies the ‘usual’ point of view (cf. p. 977 above, and paragraph 5.4) on the choice of the cells; namely, that the cells are chosen in correspondence to our finite observation capabilities. Observability of the macroscopic variables is not sufficient for the success of the repeated randomness assumption. It is conceivable (and occurs in practice) that a particular partition in terms of observable quantities does not lead to a Markov process. In that case, the choice of observable variables is simply inadequate and has to be extended with other (unobservable) quantities until we (hopefully) obtain an exhaustive set, i.e. a set of variables for which the evolution can be described autonomously. An example is the spin-echo experiment: the (observable) total magnetization of the system does not provide a suitable coarse-grained description. For further discussion of this theme, see: [Blatt, 1959; Ridderbos and Redhead, 1998; Lavis, 2004; Balian, 2005].

Apart from the unsolved problem for which partition the repeated randomness assumption is to be applied, other objections have been raised against the repeated randomness assumption. Van Kampen actually gives us not much more than the advice to accept the repeated randomness assumption bravely, not to be distracted by its dubious status, and firmly keep our eyes on its success. For authors as [Sklar, 1993], who refers to the assumption as a “rerandomization posit”, this puts the problem on its head. They request a justification of the assumption that would *explain* the success of the approach. (Indeed, even [van Kampen, 1981, p. 80] describes this success as a “miraculous fact!”). Such a request, of course, will not be satisfied by a justification that relies on its success. (But that does not mean, in my opinion, that it is an invalid form of justification.)

Another point that seems repugnant to many authors, is that the repeated coarse-graining operations appear to be added ‘by hand’, in deviation from the true dynamical evolution provided by U_t . The increase of entropy and the approach to equilibrium would thus apparently be a consequence of the fact that *we* shake up the probability distribution repeatedly in order to wash away all information about the past, while refusing a dynamical explanation for this procedure. [Redhead, 1995, p. 31] describes this procedure as “one of the most deceitful artifices I have ever come across in theoretical physics” (see also [Blatt, 1959] [Sklar, 1993] and [Callender, 1999] for similar objections).

One might ask whether the contrast between the repeated randomness assumption and the dynamical evolution need be so bleak as Van Kampen and his critics argue. After all, as we have seen in paragraph 6.2, there are dynamical systems so high in the ergodic hierarchy that they possess the Bernoulli property for some

partition of phase space (cf. paragraph 6.2). Since the Markov property is weaker than the Bernoulli property, one may infer there are also dynamical systems whose coarse grained evolutions define a homogeneous Markov process.⁷³ Thus one might be tempted to argue that the Markov property, or the repeated randomness assumption proposed to motivate it, need not require a miraculous intervention from an external ‘hand’ that throws information away; a sufficiently complex deterministic dynamics on the microscopic phase space of the system might do the job all by itself. However, the properties distinguished in the ergodic hierarchy all rely on a given measure-preserving evolution. Thus, while some dynamical systems may have the Markov property, they only give rise to *stationary* Markov processes. Its measure-preserving dynamics still implies that the Gibbs entropy remains constant. Thus, the result (208) can only be obtained in the case when all inequality signs reduce to equalities. To obtain a non-trivial form of coarse-graining, we should indeed suspend the measure-preserving dynamics.

In conclusion,!!! although the choice of a privileged partition remains an unsolved problem, there need not be a conflict between the repeated randomness assumption and the deterministic character of the dynamics at the microscopic level. However, whether the assumption (206) might actually hold for Hamiltonian systems interesting for statistical mechanics is, as far as I know, still open.

Interventionism or ‘open systems’

Another approach to stochastic dynamics is by reference to open systems. The idea is here that the system in continual interaction with the environment, and that this is responsible for the approach to equilibrium.

Indeed, it cannot be denied that in concrete systems isolation is an unrealistic idealization. The actual effect of interaction with the environment on the microscopic evolution can be enormous. A proverbial example, going back to [Borel, 1914], estimates the gravitational effect caused by displacing one gram of matter on Sirius by one centimeter on the microscopic evolution of an earthly cylinder of gas. Under normal conditions, the effect is so large, that, roughly and for a typical molecule in the gas, it may be decisive for whether or not this molecule will hit another given molecule after about 50 intermediary collisions. That is to say: microscopic dynamical evolutions corresponding to the displaced and the undisplaced matter on Sirius start to diverge considerably after a time of about 10^{-6} sec. In other words, the mechanical evolution of such a system is so extremely sensitive for disturbances of the initial state that even the most minute changes in the state of the environment can be responsible for large changes in the microscopic trajectory. But we cannot control the state of environment. Is it possible to regard irreversible behaviour as the result of such uncontrollable disturbances

⁷³Strictly speaking this is true only for discrete dynamical systems. For continuous time, e.g. for Hamiltonian dynamics, the Markov property can only be obtained by adding a time smoothing procedure to the repeated randomness assumption [Emch, 1965],[Emch and Liu, 2001, pp. 484–486].

from outside?⁷⁴

Let (x, y) be the state of a total system, where, as before, $x \in \Gamma^{(s)}$ represents the state of the object system and $y \in \Gamma^{(e)}$ that of the environment. We assume that the total system is governed by a Hamiltonian of the form

$$(209) \quad H_{\text{tot}}(x, y) = H_{(s)} + H_{(e)} + \lambda H_{\text{int}}(x, y),$$

so that the probability density of the ensemble of total systems evolves as

$$(210) \quad \rho_t(x, y) = U_t^* \rho_0(x, y) = \rho(U_{-t}(x, y))$$

i.e., a time-symmetric, deterministic and measure-preserving evolution.

At each time, we may define marginal distributions for both system and environment:

$$(211) \quad \rho_t^{(s)}(x) = \int dy \rho_t(x, y),$$

$$(212) \quad \rho_t^{(e)}(y) = \int dx \rho_t(x, y).$$

We are, of course, mostly interested in the object system, i.e. in (211). Assume further that at time $t = 0$ the total density factorizes:

$$(213) \quad \rho_0(x, y) = \rho_0^{(s)}(x) \rho_0^{(e)}(y).$$

What can we say about the evolution of $\rho_t^{(s)}(x)$? Does it form a Markov process, and does it show increase of entropy?

An immediate result (see e.g. [Penrose and Percival, 1962]) is this:

$$(214) \quad \sigma[\rho_t^{(s)}] + \sigma[\rho_t^{(e)}] \geq \sigma[\rho_0^{(s)}] + \sigma[\rho_0^{(e)}],$$

where σ denotes the Gibbs fine-grained entropy (90). This result follows from the fact that $\sigma[\rho_t]$ is conserved and that the entropy of a joint probability distribution is always smaller than or equal to the sum of the entropies of their marginals; with equality if the joint distribution factorizes. This gives a form of entropy change for the total system, but it is not sufficient to conclude that the object system itself will evolve towards equilibrium, or even that its entropy will be monotonically increasing. (Notice that (214) holds for $t \leq 0$ too.)

Actually, this is obviously not to be expected. There are interactions with an environment that may lead the system away from equilibrium. We shall have to make additional assumptions about the situation. A more or less usual set of assumptions is:

- (a). The environment is very large (or even infinite); i.e.: the dimension of $\Gamma^{(e)}$ is much larger than that of $\Gamma^{(s)}$, and $H_{(s)} \ll H_{(e)}$.

⁷⁴Note that the term ‘open system’ is employed here for a system in (weak) interaction with its environment. This should be distinguished from the notion of ‘open system’ in other branches of physics where it denotes a system that can exchange particles with its environment.

- (b). The coupling between the system and the environment is weak, i.e. λ is very small.
- (c). The environment is initially in thermal equilibrium, e.g., $\rho^{(e)}(y)$ is canonical:

$$(215) \quad \rho_0^{(e)} = \frac{1}{Z(\beta)} e^{-\beta H^{(e)}}$$

- (d). One considers time scales only that are long with respect to the relaxation times of the environment, but short with respect to the Poincaré recurrence time of the total system.

Even then, it is a major task to obtain a master equation for the evolution of the marginal state (211) of the system, or to show that its evolution is generated by a semigroup, which would guarantee that this forms a Markov process (under the proviso of footnote 71). Many specific models have been studied in detail (cf. [Spohn, 1980]). General theorems were obtained (although mostly in a quantum mechanical setting) by [Davies, 1974; Davies, 1976a; Lindblad, 1976; Gorini et al., 1976]. But there is a similarity to the earlier approach: it seems that, here too, an analogue of ‘repeated randomness’ must be introduced. [Mehra and Sudarshan, 1972; van Kampen, 1994; Maes and Netočný, 2003].

At the risk of oversimplifying and misrepresenting the results obtained in this analysis, I believe they can be summarized as showing that, in the so-called ‘weak coupling’ limit, or some similar limiting procedure, the time development of (211) can be modeled as

$$(216) \quad \rho_t^{(s)}(x) = T_t \rho^{(s)}(x) \quad t \geq 0,$$

where the operators T_t form a semigroup, while the environment remains in its steady equilibrium state:

$$(217) \quad \rho_t^{(e)}(y) = \rho_0^{(e)}(y) \quad t \geq 0.$$

The establishment of these results would also allow one to infer, from (214), the monotonic non-decrease of entropy of the system.

To assess these findings, it is convenient to define, for a fixed choice of $\rho_0^{(e)}$ the following linear map on probability distributions of the total system:

$$(218) \quad \mathcal{TR} : \rho(x, y) \mapsto \mathcal{TR}\rho(x, y) = \int \rho(x, y) dy \cdot \rho_0(y)$$

This map removes the correlation between the system and the environment, and projects the marginal distribution of the environment back to its original equilibrium form.

Now, it is not difficult to see that the Chapman-Kolmogorov equation (which is equivalent to the semigroup property) can be expressed as

$$(219) \quad \mathcal{TR}U_{t+t'}^* = \mathcal{TR}U_{t'}^* \mathcal{TR}U_t^* \quad \text{for all } t, t' \geq 0$$

which is analogous to (206).

There is thus a strong formal analogy between the coarse-graining and the open-systems approaches. Indeed, the variables of the environment play a role comparable to the internal coordinates of a cell in the coarse graining approach. The exact microscopic information about the past is here translated into the form of correlations with the environment. This information is now removed by assuming that at later times, effectively, the state may be replaced by a product of the form (213), neglecting the back-action of the system on the environment. The mappings \mathcal{CG} and \mathcal{TR} are both linear and idempotent mappings, that can be regarded as special cases of the projection operator techniques of Nakajima and Zwanzig, which allows for a more systematical and abstract elaboration, sometimes called *subdynamics*.

Some proponents of the open systems approach, (e.g. [Morrison, 1966; Redhead, 1995]), argue that in contrast to the coarse-graining approach, the present procedure is ‘objective’. Presumably, this means that there is supposed to be a fact of the matter about whether the correlations are indeed ‘exported to the environment’. However, the analogy between both approaches makes one suspect that any problem for the coarse-graining approach is translated into an analogous problem of the open systems approach. Indeed, the problem of finding a privileged partition that we discussed in the previous paragraph is mirrored here by the question where one should place the division between the ‘system’ and ‘environment’. There is no doubt that in practical applications this choice is also arbitrary.

7.6 *Can the Markov property explain irreversible behaviour?*

Ad (v). Finally, I turn to what may well be the most controversial and surprising issue: is the Markov property, or the repeated randomness assumption offered to motivate it, responsible for the derivation of time-reversal non-invariant results?

We have seen that every non-invertible homogeneous Markov process displays “irreversible behaviour” in the sense that different initial probability distributions will tend to become more alike in the course of time. Under certain technical conditions, one can obtain stronger results, e.g. an approach to a unique equilibrium state, monotonic non-decrease of absolute entropy, etc. All these results seem to be clearly time-asymmetric. And yet we have also seen that the Markov property is explicitly time symmetric. How can these be reconciled?

To start off, it may be noted that it has often been affirmed that the Markov property is the key towards obtaining time-asymmetric results. For example, Penrose writes:

“...the behaviour of systems that are far from equilibrium is not symmetric under time reversal: for example: heat always flows from a hotter to a colder body, never from a colder to a hotter. If this behaviour could be derived from the symmetric laws of dynamics alone there would, indeed, be a paradox; we must therefore acknowledge the fact that some additional postulate, non-dynamical in character and

asymmetric under time reversal must be adjoined to the symmetric laws of dynamics before the theory can become rich enough to explain non-equilibrium behaviour. In the present theory, this additional postulate is the Markov postulate” (Penrose 1970, p. 41).

In the previous paragraph, we have already questioned the claim expressed here that the Markov property is “non-dynamical”. But now we are interested in the question whether postulating the Markov property would be asymmetric under time-reversal. Many similar statements, e.g. that the repeated randomness assumption is “the additional element by which statistical mechanics has to be supplemented in order to obtain irreversible equations” [van Kampen, 1962, p. 182], or that the non-invertibility of a Markov process provides the origin of thermodynamic behaviour [Mackey, 1992] can be found in the works of advocates of this approach.

But how can this be, given that the Markov property is explicitly time-symmetric? In order to probe this problem, consider another question. How does a given probability distribution $P(y, 0)$ evolve for negative times? So, starting again from (170), let us now take $t \leq 0$. We still have:

$$(220) \quad P(y, t) = \sum_{y'} P(y, t, |y', 0)P(y', 0).$$

These conditional probabilities $P(y, t, |y', 0)$ satisfy the ‘time-reversed’ Markov property (174), that says that extra specification of later values is irrelevant for the retrodiction of earlier values. As a consequence, we get for $t \leq t' \leq t'', 0$:

$$(221) \quad P(y, t|y'', t'') = \sum_{y'} P(y, t|y', t')P(y', t'|y'', t'')$$

i.e., a time-reversed analogue of the Chapman-Kolmogorov equation.

We may thus also consider these conditional probabilities for negative times as backward evolution operators. If we could assume their invariance under time translation, i.e. that they depend only on the difference $\tau = t - t'$, we could write

$$(222) \quad S_\tau(y|y') := P(y, t|y, t') \quad \text{with } \tau = t - t' \leq 0,$$

and obtain a second semigroup of operators S_τ , obeying

$$(223) \quad S_{\tau+\tau'} = S_\tau \circ S_{\tau'} \quad \tau, \tau' \leq 0$$

that generate stochastic evolutions towards the past.

Further, these backward conditional probabilities are connected to the forward conditional probabilities by means of Bayes’ theorem:

$$(224) \quad P_{(1|1)}(y, t|y', t') = \frac{P_{(1|1)}(y', t'|y, t)P(y, t)}{P(y', t')};$$

and if the process, as before, is homogeneous this becomes

$$(225) \quad P_{(1|1)}(y, t|y', t') = \frac{T_{-\tau}(y'|y)P_t(y)}{P_{t'}(y')} \quad ; \quad \tau = t - t' < 0.$$

The matrix $P_{(1|1)}(y, t|y', t')$ always gives for $t < t'$ the correct ‘inversion’ of T_t . That is to say:

$$(226) \quad \sum_{y'} P(y, t|y', t')(T_{t'-t}P_t)(y') = P_t(y)$$

Note firstly that (225) is *not* the matrix-inverse of T_t ! Indeed, the right-hand side of (225) depends on P_t and $P_{t'}$ as well as T . Even if the matrix-inverse $T^{(inv)}$ does not exist, or is not a bona fide stochastic matrix, the evolution towards the past is governed by the Bayesian inversion, i.e. by the transition probabilities (225).

Note also that if the forward transition probabilities are homogeneous, this is not necessarily so for the backward transition probabilities. For example, if in (225) one translates both t and t' by δ , one finds

$$P(y, t + \delta|y', t' + \delta) = \frac{T_{-\tau}(y'|y)P(y, t + \delta)}{P(y', t' + \delta)}.$$

Here, the right-hand side generally still depends on δ . In the special case that the initial distribution is itself stationary, the backward transition probabilities are homogeneous whenever the forward ones are. If $P(y, t)$ is not stationary, we might still reach the same conclusion, as long as the non-stationarity is restricted to those elements y or y' of \mathcal{Y} for which $T_t(y|y') = 0$ for all t . Otherwise, the two notions become logically independent.

This gives rise to an unexpected new problem. Usually, an assumption of homogeneity (or time translation invariance) is seen as philosophically innocuous, as compared to time reversal invariance. But here we see that assuming time translation invariance for a system of *forward* transition probabilities is not equivalent to assuming the same invariance for the *backward* transition probabilities. If one believes that one of the two is obvious, how will one go about explaining the failure of the other? And how would one explain the preference for which one of the two is obvious, without falling into the “double standards” accusation of the kind raised by [Price, 1996]?

But what about entropy increase? We have seen before that for every non-invertible Markov process the relative entropy of the distribution P with respect to the equilibrium distribution increases, and that the distribution evolves towards equilibrium. (Homogeneity of the process is not needed for this conclusion.) But the backward evolution operators form a Markov process too, for which exactly the same holds. This seems paradoxical. If $T_t P_0 = P_t$, we also have $P_t = S_{-t} P_0$. The entropy of P_t can hardly be both higher and lower than that of P_0 ! An example may clarify the resolution of this apparent problem: namely, the stationary solutions of S are not the same as the stationary solutions of T !

Example Consider a Markov chain with $\mathcal{Y} = \{1, 2\}$ and let

$$(227) \quad T = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

Choose an initial distribution $P_0 = \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix}$. After one step we already get

$$(228) \quad TP = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$$

which is also the (unique) stationary distribution P^* . The backward transition probabilities are given by Bayes' theorem, and one finds easily:

$$(229) \quad S = \begin{pmatrix} \alpha & \alpha \\ 1 - \alpha & 1 - \alpha \end{pmatrix}.$$

The stationary distribution for this transition probability is

$$(230) \quad \tilde{P}^* = \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix}.$$

That is to say: for the forward evolution operator the transition

$$(231) \quad \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix} \xrightarrow{T} \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}$$

is one for which a non-stationary initial distribution evolves towards a stationary one. The relative entropy increases: $H(P_0, P^*) \leq H(TP, P^*)$. But for the backward evolution, similarly:

$$(232) \quad \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} \xrightarrow{S} \begin{pmatrix} \alpha \\ 1 - \alpha \end{pmatrix}$$

represents an evolution from a non-stationary initial distribution to the stationary distribution \tilde{P}^* and, here too, relative entropy increases: $H(P_1, \tilde{P}^*) \leq H(P_0, \tilde{P}^*)$.

The illusion that non-invertible Markov processes possess a built-in time-asymmetry is (at least partly) due to the habit of regarding T_τ as a fixed evolution operator on an independently chosen distribution P_0 . Such a view is of course very familiar in other problems in physics, where deterministic evolution operators generally *do* form a group and may be used, at our heart's desire, for positive and negative times.

Indeed, the fact that these operators in general have no inverse might seem to reflect the idea that Markov processes have no memory and 'lose information' along the way and that is the cause of the irreversible behaviour, embodied in the time-asymmetric master equation, increase of relative or absolute entropy or approach to equilibrium. But actually, every Markov process has apart from a system of forward, also a system of backward transition probabilities, that again forms a semigroup (when they are homogeneous). If we had considered *them* as given we would get all conclusions we obtained before, but now for negative times.

I conclude that irreversible behaviour is not built into the Markov property, or in the non-invertibility of the transition probabilities, (or in the repeated randomness assumption⁷⁵, or in the Master equation or in the semigroup property). Rather the appearance of irreversible behaviour is due to the choice to rely on the forward transition probabilities, and not the backward. A similar conclusion has been reached before [Edens, 2001] in the context of proposals of Prigogine and his coworkers. My main point here is that the same verdict also holds for more ‘mainstream’ approaches as coarse-graining or open systems.

7.7 Reversibility of stochastic processes

In order not to end this chapter on a destructive note, let me emphasize that I do not claim that the derivation of irreversible behaviour in stochastic dynamics is impossible. Instead, the claim is that motivations for desirable properties of the forward transition probabilities are not enough; one ought also show that these properties are lacking for the backward transitions.

In order to set up the problem of irreversibility in this approach to non-equilibrium statistical mechanics for a more systematic discussion, one first ought to provide a reasonable definition for what it means for a stochastic process to be (ir)reversible; a definition that would capture the intuitions behind its original background in Hamiltonian statistical mechanics.

One general definition that seems to be common (cf. [Kelly, 1979 p. 5]) is to call a stochastic process reversible iff, for all n and t_1, \dots, t_n and τ :

$$(233) \quad P_{(n)}(y_1, t_1; \dots; y_n, t_n) = P_{(n)}(y_1, \tau - t_n; \dots; y_n, \tau - t_n).$$

See [Grimmett and Stirzaker, 1982, p. 219] for a similar definition restricted to Markov processes) The immediate consequence of this definition is that a stochastic process can only be reversible if the single-time probability $P_{(1)}(y, t)$ is stationary, i.e. in statistical equilibrium. Indeed, this definition seems to make the whole problem of reconciling irreversible behaviour with reversibility disappear. As [Kelly, 1979, p. 19] notes in a discussion of the Ehrenfest model: “there is no conflict between reversibility and the phenomenon of increasing entropy — reversibility is a property of the model in equilibrium and increasing entropy is a property of the approach to equilibrium”

But clearly, this view trivializes the problem, and therefore it is not the appropriate definition for non-equilibrium statistical mechanics. Recall that the Ehrenfest dog flea model (§7.2) was originally proposed in an attempt of showing how a tendency of approaching equilibrium from a initial non-equilibrium distribution (e.g.

⁷⁵In recent work, van Kampen acknowledges that the repeated randomness assumption by itself does not lead to irreversibility: “This repeated randomness assumption [...] breaks the time symmetry by explicitly postulating the randomization *at the beginning* of the time interval Δt . There is no logical justification for this assumption other than that it is the only thing one can do and that it works. If one assumes randomness at the end of each Δt coefficients for diffusion, viscosity, etc. appear with the wrong sign; if one assumes randomness at the midpoint no irreversibility appears” [van Kampen, 2002, p.475, original emphasis].

a probability distribution that gives probability 1 to the state that all fleas are located on the same dog) could be reconciled with a stochastic yet time-symmetric dynamics.

If one wants to separate considerations about initial conditions from dynamical considerations at all, one would like to provide a notion of (ir)reversibility that is associated with the stochastic dynamics alone, independent of the initial distribution is stationary.

It seems that an alternative definition which would fulfill this intuition is to say that a stochastic process is reversible if, for all y and y' and $t' > t$,

$$(234) \quad P_{(1|1)}(y, t|y', t') = P_{(1|1)}(y, t'|y', t).$$

In this case we cannot conclude that the process must be stationary, and indeed, the Ehrenfest model would be an example of a reversible stochastic process. I believe this definition captures the intuition that if at some time state y' obtains, the conditional probability of the state one time-step earlier being y is equal to that of the state one time-step later being y .

According to this proposal, the aim of finding the “origin” of irreversible behaviour or “time’s arrow”, etc. in stochastic dynamics must then lie in finding and motivating conditions under which the forward transition probabilities are different from the backwards transition probabilities, in the sense of a violation of (234). Otherwise, irreversible behaviour would essentially be a consequence of the assumptions about initial conditions, a result that would not be different in principle from conclusions obtainable from Hamiltonian dynamics.

BIBLIOGRAPHY

- [Albert, 2000] D. Z. Albert. *Time and Chance* Cambridge, Mass.: Harvard University Press, 2000.
- [Alberti and Uhlmann, 1982] P. M. Alberti and A. Uhlmann. *Stochasticity and partial order. Doubly stochastic maps and unitary mixing*. Dordrecht: Reidel, 1982.
- [Balescu, 1997] R. Balescu. *Statistical dynamics*. London: Imperial College Press, 1997.
- [Balian, 2005] R. Balian. Information in statistical physics. *Studies In History and Philosophy of Modern Physics*, 36, 323–353, 2005.
- [Batterman, 1990] R. W. Batterman. Irreversibility and statistical mechanics: A new approach? *Philosophy of Science*, 57, 395–419, 1990.
- [Batterman, 1991] R. W. Batterman. Randomness and probability in dynamical theories: On the proposals of the Prigogine school. *Philosophy of Science*, 58, 241–263, 1991.
- [Batterman, 2002] R. W. Batterman. *The devil in the details: asymptotic reasoning in explanation, reduction, and emergence* Oxford: Oxford University Press, 2002.
- [Batterman, 2005] R. W. Batterman. Critical phenomena and breaking drops: Infinite idealizations in physics. *Studies in History and Philosophy of Modern Physics*, 36, 225–244, 2005.
- [Becker, 1967] R. Becker. *Theory of heat* New York: Springer, 1967.
- [Belot, 2006] G. Belot. This volume, chapter 2, 2006.
- [Bennett, 2003] C. H. Bennett. Notes on Landauer’s principle, reversible computation, and Maxwell’s demon. *Studies In History and Philosophy of Modern Physics*, 34, 501–510, 2003.
- [Berkovitz et al., 2006] Berkovitz, J., R. Frigg and F. Kronz. The ergodic hierarchy and Hamiltonian chaos. *Studies in History and Philosophy of Modern Physics*, 2006.
- [Bertrand, 1889] J. Bertrand. *Calcul des Probabilités*. Paris: Gauthier-Villars. Reissued New York: Chelsea, 1889. no date.

- [Bierhalter, 1994] G. Bierhalter. Von L. Boltzmann bis J.J. Thomson: die Versuche einer mechanischen Grundlegung der Thermodynamik (1866-1890). *Archive for History of Exact Sciences*, 44, 25–75, 1994.
- [Birkhoff, 1931] G. D. Birkhoff. Proof of the ergodic theorem. *Proceedings of the National Academy of Sciences of the United States of America*, 17, 656–660, 1931.
- [Bishop, 2004] R. C. Bishop. Nonequilibrium statistical mechanics Brussels-Austin style *Studies In History and Philosophy of Modern Physics*, 35, 1–30, 2004.
- [Blatt, 1959] J. M. Blatt. An alternative approach to the ergodic problem. *Progress in Theoretical Physics*, 22, 745–756, 1959.
- [Boltzmann, 1866] L. Boltzmann. Über die Mechanische Bedeutung des Zweiten Hauptsatzes der Wärmetheorie. *Wiener Berichte*, 53, 195–220 1866; in [Boltzmann, 1909] Vol. I, paper 2.
- [Boltzmann, 1868a] L. Boltzmann. Studien über das Gleichgewicht der lebendigen Kraft zwischen bewegten materiellen Punkten. *Wiener Berichte*, 58, 517–560, 1868; in [Boltzmann, 1909] Vol. I, paper 5.
- [Boltzmann, 1868b] L. Boltzmann. Lösung eines mechanischen Problems. *Wiener Berichte*, 58, 1035–1044, 1868. In [Boltzmann, 1909] Vol. I, paper 6.
- [Boltzmann, 1871a] L. Boltzmann. Über das Wärme Gleichgewicht zwischen mehratomigen Gasmolekülen. *Wiener Berichte*, 63, 397–418, 1871. In [Boltzmann, 1909] Vol. I, paper 18.
- [Boltzmann, 1871b] L. Boltzmann. Einige allgemeine Sätze über Wärme Gleichgewicht. *Wiener Berichte*, 63, 679–711, 1871. In [Boltzmann, 1909] Vol. I, paper 19.
- [Boltzmann, 1871c] L. Boltzmann. Analytischer Beweis des zweiten Hauptsatzes der mechanischen Wärmetheorie aus den Sätzen über das Gleichgewicht der lebendigen Kraft. *Wiener Berichte*, 63, 712–732, 1871. In [Boltzmann, 1909] Vol. I, paper 20.
- [Boltzmann, 1872] L. Boltzmann. Weitere Studien über das Wärme Gleichgewicht unter Gasmolekülen. *Wiener Berichte*, 66, 275–370, 1872. In [Boltzmann, 1909] Vol. I, paper 23.
- [Boltzmann, 1877a] L. Boltzmann. . Bemerkungen über einige Probleme der mechanische Wärmetheorie *Wiener Berichte*, 75, 62–100, 1877. In [Boltzmann, 1909] Vol. II, paper 39.
- [Boltzmann, 1877b] L. Boltzmann. Über die Beziehung zwisschen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung resp. dem Sätzen über das Wärme Gleichgewicht *Wiener Berichte*, 76, 373–435, 1877. In [Boltzmann, 1909] Vol. II, paper 42.
- [Boltzmann, 1878] L. Boltzmann. Weitere Bemerkungen über einige Probleme der mechanischen Wärmetheorie. *Wiener Berichte*, 78, 7–46, 1878. In [Boltzmann, 1909] Vol. II paper 44.
- [Boltzmann, 1881a] L. Boltzmann. Über einige das Wärme Gleichgewicht betreffende Sätze. *Wiener Berichte*, 84, 136–145, 1881. In [Boltzmann, 1909] Vol. II paper 62.
- [Boltzmann, 1881b] L. Boltzmann. Referat über die Abhandlung von J.C. Maxwell: “Über Boltzmann’s Theorem betreffend die mittlere verteilung der lebendige Kraft in einem System materieller Punkte”. *Wiedemann’s Annalen Beiblätter*, 5, 403–417, 1881. In [Boltzmann, 1909] Vol. II paper 63.
- [Boltzmann, 1884] L. Boltzmann. Über die Eigenschaften Monocyklischer und andere damit verwandter Systeme. *Crelle’s Journal für die reine und angewandte Mathematik*, 98, 68–94 1884 and 1885. In [Boltzmann, 1909] Vol III, paper 73.
- [Boltzmann, 1887a] L. Boltzmann. Ueber die mechanischen Analogien des zweiten Hauptsatzes der Thermodynamik. *Journal für die reine und angewandte Mathematik*, 100, 201–212, 1887. Also in [Boltzmann, 1909], Vol. III, paper 78.
- [Boltzmann, 1887b] L. Boltzmann. . Neuer Beweis zweier Sätze über das Wärme Gleichgewicht unter mehratomigen Gasmolekülen. *Wiener Berichte* 95, 153–164, 1887 in [Boltzmann, 1909] Vol. III, paper 83.
- [Boltzmann, 1887c] L. Boltzmann. Über einige Fragen der Kinetische Gastheorie. *Wiener Berichte* 96, 891–918, 1887. In [Boltzmann, 1909] Vol. III, paper 86.
- [Boltzmann, 1892] L. Boltzmann. III. Teil der Studien über Gleichgewicht der lebendigen Kraft *Münch. Ber.* 22, 329–358, 1892. In [Boltzmann, 1909] Vol. III, paper 97.
- [Boltzmann, 1894a] L. Boltzmann and C.H. Bryan. Über die mechanische Analogie des Wärme Gleichgewichtes zweier sich berührende Körper. *Wiener Berichte*, 103, 1125–1134, 1894. In [Boltzmann, 1909] Vol. III, paper 107.
- [Boltzmann, 1894b] L. Boltzmann. . On the application of the determinantal relation to the kinetic theory of gases. Appendix C to an article by G.H. Bryan on thermodynamics. *Reports of the British Association for the Advancement of Science*, pp. 102–106, 1894. In [Boltzmann, 1909] Vol. III, paper 108.

- [Boltzmann, 1895] L. Boltzmann. On certain questions in the theory of gases. *Nature* **51**, 413–415, 1895. In [Boltzmann, 1909], Vol. III, paper 112.
- [Boltzmann, 1895b] L. Boltzmann. On the minimum theorem in the theory of gases. *Nature* **52**, 221, 1895. Also in [Boltzmann, 1909] Vol. III, paper 114.
- [Boltzmann, 1896] L. Boltzmann. *Vorlesungen über Gastheorie* Vol I. Leipzig, J.A. Barth, 1896. Translated, together with [Boltzmann, 1898] by S.G. Brush, *Lecture on Gas Theory* Berkeley: University of California Press, 1964.
- [Boltzmann, 1896b] L. Boltzmann. Entgegnung an die wärmetheoretischen Betrachtungen des Hrn. E., Zermelo. *Wiedemann's Annalen*, *57*, 772–784, 1896. In [Boltzmann, 1909] Vol. III, paper 119.
- [Boltzmann, 1897a] L. Boltzmann. Zu Hrn Zermelos Abhandlung “Über die mechanische Erklärung irreversibler Vorgänge”. *Wiedemann's Annalen*, *60* 392–398, 1897. in [Boltzmann, 1909], Vol. III paper 120.
- [Boltzmann, 1897b] L. Boltzmann. Über einige meiner weniger bekannte Abhandlungen über Gastheorie und deren Verhältnis zu derselben *Verhandlungen des 69. Versammlung Deutscher Naturforscher und Ärzte*, Braunschweig 1897, pp. 19–26. *Jahresberichte der Deutsche Mathematiker-Vereinigung*, **6**, (1899), 130–138. Also in [Boltzmann, 1909] Vol. III, paper 123.
- [Boltzmann, 1898] L. Boltzmann. *Vorlesungen über Gastheorie* Vol II. Leipzig, J.A. Barth, 1898. Translated, together with [Boltzmann, 1896] by S.G. Brush, *Lecture on Gas Theory* Berkeley: University of California Press, 1964.
- [Boltzmann, 1898] L. Boltzmann. Über die sogenannte *H*-Kurve. *Mathematische Annalen* **50**, 325–332, 1898. In [Boltzmann, 1909] vol. III, paper 127.
- [Boltzmann and Nabl, 1904] L. Boltzmann and J. Nabl. Kinetisch theorie der Materie. *Encyclopädie der Mathematischen Wissenschaften*, Vol V-1, pp. 493–557, 1904.
- [Boltzmann, 1905] L. Boltzmann. *Populäre Schriften*. Leipzig: J.A. Barth, 1905. Re-issued Braunschweig: F. Vieweg, 1979.
- [Boltzmann, 1909] L. Boltzmann. *Wissenschaftliche Abhandlungen* Vol. I, II, and III. F. Hasenöhrl (ed.) Leipzig, 1909. Reissued New York: Chelsea, 1969.
- [Borel, 1914] E. Borel. *Le Hasard*. Paris: Alcan, 1914.
- [Boyling, 1972] J. B. Boyling. An axiomatic approach to classical thermodynamics. *Proceedings of the Royal Society of London*, *329*, 35–71, 1972.
- [Bricmont, 1995] J. Bricmont. Science of chaos or chaos in science? *Physicalia*, *17*, 159–208; also in P.R. Gross, N.Levitt and M.W.Lewis (Eds.), *The flight from science and reason*, New York: New York Academy of Sciences, pp. 131–176, 1995.
- [Brown and Uffink, 2001] H. Brown and J. Uffink. (2001). The origins of time-asymmetry in thermodynamics: the minus first law. *Studies in History and Philosophy of Modern Physics*, *32*, 525–538, 2001.
- [Brush, 1966] S. G. Brush. *Kinetic theory* Vol. 1 and 2. Oxford, 1966: Pergamon.
- [Brush, 1976] S. G. Brush. *The Kind of motion we call heat*. Amsterdam: North Holland, 1976.
- [Brush, 1977] S. G. Brush. Statistical mechanics and the philosophy of science: some historical notes. In F. Supper & P.D. Asquith (Eds.) *PSA 1976 Proceedings of the Biennial meeting of the Philosophy of Science Association 1976, Vol 2*, East Lansing, Mich.: Philosophy of Science Association, pp. 551–584, 1976.
- [Brush, 1999] S. G. Brush. Gadflies and geniuses in the history of gas theory. *Synthese* *119*, 11–43. Also in [Brush, 2003], pp. 421–450, 1999.
- [Brush, 2003] S. G. Brush. *The kinetic theory of gases*. London: Imperial College Press, 2003.
- [Bryan, 1894] G. H. Bryan. Letter to the editor. *Nature*, *51*, 175, 1894.
- [Bryan, 1894] G. H. Bryan. The Assumption in Boltzmann's Minimum Theorem, *Nature*, *52*, 29–30, 1894.
- [Burbury, 1894a] S. H. Burbury. Boltzmann's minimum theorem. *Nature*, *51*, 78–79, 1894.
- [Burbury, 1894b] S. H. Burbury. The kinetic theory of gases. *Nature*, *51* 175–176, 1894.
- [Butterfield, 2006] J. N. Butterfield. This volume, chapter 1, 2006.
- [Callender, 1999] C. Callender. Reducing thermodynamics to statistical mechanics: the case of entropy. *Journal of Philosophy*, *96*, 348–373, 1999.
- [Callender, 2001] C. Callender. Taking Thermodynamics Too Seriously. *Studies In History and Philosophy of Modern Physics*, *32*, 539–553, 2001.
- [Callender, 2004] C. Callender. Measures, explanations and the past: Should ‘special’ initial conditions be explained? *British Journal for Philosophy of Science*, *55*, 195–217, 2004.

- [Campisi, 2005] M. Campisi. On the mechanical foundations of thermodynamics: The generalized Helmholtz theorem *Studies In History and Philosophy of Modern Physics*, 36, 275–290, 2005.
- [Casetti et al, 2003] L. Casetti, M. Pettini and E. G. D. Cohen. Phase Transitions and Topology Changes in Configuration Space. *Journal of Statistical Physics*, 111, 1091–1123, 2003.
- [Carathéodory, 1909] C. Carathéodory. Untersuchungen über die Grundlagen der Thermodynamik, *Mathematische Annalen*, 67, 355–386, 1909. Translation by J. Kestin, “Investigation into the foundations of thermodynamics” in [Kestin, 1976], pp. 229–256. This translation is not quite accurate.
- [Carnot, 1824] S. Carnot. *Réflexions sur la puissance motrice du feu*. Paris: Bachelier, 1824. Re-edited and translated in E. Mendoza (Ed.) *Reflections on the motive power of fire*, New York: Dover, 1960.
- [Cercignani, 1998] C. Cercignani. *Ludwig Boltzmann, the man who trusted atoms* Oxford: Oxford University Press, 1998.
- [Cercignani et al., 1994] C. Cercignani, R. Illner, and M. Pulvirenti. *The mathematical theory of dilute gases*. New York: Springer-Verlag, 1994.
- [Chang, 2003] H. Chang. Preservative realism and its discontents: revisiting caloric. *Philosophy of Science*, 70, 902–912, 2003.
- [Chang, 2004] H. Chang. *Inventing temperature. Measurement and scientific progress*. Oxford: Oxford University Press, 2004.
- [Clausius, 1857] R. Clausius Über die Art von Bewegung die wir Wärme nennen, *Poggendorff's Annalen*, 100, 253–280, 1857. English translation in [Brush, 1966], Vol.1. pp. 111–134.
- [, Clausius] 1862] Clausius R. Clausius. ‘Ueber die Anwendung des Satzes von der Aequivalenz der Verwandlungen auf die innere Arbeit’, *Vierteljahrsschrift der Züricher naturforschenden Gesellschaft*, 7, 48, 1862. Also in *Annalen der Physik und Chemie*, 116: 73–112 (1862), English translation in *Philosophical Magazine*, 24, 81–97, 201–213, also in [Kestin, 1976], pp. 133–161.
- [Clausius, 1865] R. Clausius. Über verschiedene für die Anwendung bequeme Formen der Hauptgleichungen der mechanische Wärmetheorie, *Poggendorff's Annalen*, 100, (253), 1865. English translation in [Kestin, 1976].
- [Cohen, 1996] E. G. D. Cohen. Boltzmann and statistical mechanics, cond-mat/9608054v2, 1996.
- [Compagner, 1989] A. Compagner. Thermodynamics as the continuum limit of statistical mechanics *American Journal of Physics*, 57, 106–117, 1989.
- [Cornfeld et al., 1982] I. P. Cornfeld, S. V. Fomin, and Ya. G. Sinai. *Ergodic theory* New York: Springer-Verlag, 1982.
- [Culverwell, 1894] E. Culverwell. Dr. Watson's proof of Boltzmann's theorem on permanence of distributions *Nature*, 50 617, 1894.
- [Culverwell, 1895] E. Culverwell. Boltzmann's minimum theorem *Nature*, 51 246, 1895.
- [Curd, 1982] M. Curd. Popper on the direction of time. In R. Sexl and J. Blackmore (eds.), *Ludwig Boltzmann, internationale Tagung: anlässlich des 75. Jahrestages seines Todes, 5.-8. September 1981: ausgewählte Abhandlungen*, pp. 263–303, 1982. Graz: Akademische Druck- und Verlagsanstalt.
- [Davies, 1974] E. B. Davies. Markovian master equations. *Communications in Mathematical Physics*, 39 91–110, 1974.
- [Davies, 1976a] E. B. Davies. Markovian master equations II *Mathematische Annalen*, 219 147–158, 1976.
- [Davies, 1976] E. B. Davies. *Quantum theory of open systems*. New York: Academic Press, 1976.
- [Denbigh and Denbigh, 1985] K. G. Denbigh and J. Denbigh. *Entropy in relation to incomplete knowledge*. Cambridge: Cambridge University Press, 1985.
- [Dias, 1994] P. M. C. Dias. Will someone say exactly what the *H*-theorem proves? A study of Burbury's Condition A and Maxwell's Proposition II. *Archive for History of Exact Sciences* 46, 341–366, 1994.
- [Dugas, 1959] R. Dugas. *La théorie physique au sens de Boltzmann et ses prolongements modernes*. (Neuchâtel: Griffon), 1959.
- [Earman, 2006] J. Earman. The past-hypothesis: not even false. *Studies in History and Philosophy of Modern Physics*, 2006.

- [Earman and Rédei, 1996] J. Earman and M. Rédei. Why ergodic theory does not explain the success of equilibrium statistical mechanics. *British Journal for the Philosophy of Science*, 45 63–78, 1996.
- [Earman and Norton, 1998] J. Earman and J. D. Norton. Exorcist XIV: The wrath of Maxwell's demon. Part I. From Maxwell to Szilard. *Studies in History and Philosophy of Modern Physics*, 29, 435–471, 1998.
- [Earman and Norton, 1999] J. Earman and J. D. Norton. Exorcist XIV: The wrath of Maxwell's demon. Part II. From Szilard to Landauer and beyond *Studies in History and Philosophy of Modern Physics*, 30, 1–40.
- [Edens, 2001] B. Edens. Semigroups and symmetry: an investigation of Prigogine's theories. <http://philsci-archive.pitt.edu/archive/00000436/>, 2001.
- [Ehrenfest and Ehrenfest-Afanassjewa, 1912] P. Ehrenfest and T. Ehrenfest-Afanassjewa. Begriffliche Grundlagen der Statistischen Auffassung in der Mechanik. *Enzyklopädie der Mathematischen Wissenschaften* Vol. 4. F. Klein and C. Müller (eds.). Leipzig: Teubner, pp. 3–90, 1912. English translation *The conceptual foundations of the statistical approach in mechanics*. Ithaca N.Y.: Cornell University Press, 1959.
- [Einstein, 1902] A. Einstein. Kinetische Theorie des Wärmegleichgewicht und des zweiten Hauptsatzes der Thermodynamik. *Annalen der Physik*, 9, 417–433, 1902.
- [Einstein, 1910] A. Einstein. Theorie der Opaleszenz von homogenen Flüssigkeiten und Flüssigkeitsgemischen in der Nähe des kritischen Zustandes *Annalen der Physik*, 33, 1275–1298, 1910.
- [Ellis, 1850] R. L. Ellis. Remarks on an alleged proof of the method of least squares, contained in a late number of the Edinburgh Review, in W. Walton (Ed.), *Mathematical and other Writings of R.L. Ellis*. Cambridge: Cambridge University Press, 1863, pp. 53–61.
- [Emch, 1965] G. G. Emch. On the Markov character of master-equations. *Helvetica Physica Acta*, 38, 164–171, 1965.
- [Emch, 2005] G. G. Emch. Probabilistic issues in statistical mechanics. *Studies In History and Philosophy of Modern Physics* 36, 303–322, 2005.
- [Emch, 2006] G. G. Emch. This volume, chapter 10, 2006.
- [Emch and Liu, 2001] G. G. Emch and C. Liu, *The Logic of thermostistical physics*. Berlin: Springer, 2001.
- [Farquhar, 1964] I. E. Farquhar. *Ergodic theory in statistical mechanics*, London, Interscience, 1964.
- [Feshbach, 1987] H. Feshbach. Small systems: when does thermodynamics apply? *Physics Today*, 40, 9–11, 1987.
- [Fine, 1973] T. L. Fine. *Theories of probability: An examination of foundations*. New York: Academic Press, 1973.
- [Fisher, 1964] M. E. Fisher. The free energy of a macroscopic system. *Archive for Rational Mechanics and Analysis* 17 377–410, 1964.
- [Fowler and Guggenheim, 1939] R. H. Fowler and E. Guggenheim. *Statistical Thermodynamics*. Cambridge: Cambridge University Press, 1939.
- [Fox, 1971] R. Fox. *The caloric theory of gases: from Lavoisier to Regnault*. Oxford: Clarendon Press, 1971.
- [Friedman, 1976] K. S. Friedman. A Partial vindication of ergodic theory. *Philosophy of Science*, 43, 151–162, 1976.
- [Frigg, 2004] R. Frigg. In what sense is the Kolmogorov-Sinai entropy a measure for chaotic behaviour? — bridging the gap between dynamical systems theory and communication theory. *British Journal for the Philosophy of Science*, 55, 411–434, 2004.
- [Galavotti, 2004] M. C. Galavotti. *A philosophical introduction to probability*. Chicago: University of Chicago Press, 2004.
- [Gallavotti, 1994] G. Gallavotti. Ergodicity, ensembles, irreversibility in Boltzmann and beyond <http://arXiv:chao-dyn/9403004>. *Journal of Statistical Physics*, 78, 1571–1589, 1994.
- [Gallavotti, 1999] G. Gallavotti. *Statistical mechanics: A short treatise*. Berlin: Springer, 1999.
- [Gantmacher, 1959] F. R. Gantmacher. *Matrizenrechnung* Vol 2. Berlin: Deutscher Verlag der Wissenschaften, 1959.
- [Garber et al., 1986] E. Garber, S. G. Brush, and C. W. F. Everitt, (Eds.). *Maxwell on molecules and gases* Cambridge Mass.: MIT Press, 1986.
- [Garber et al., 1995] E. Garber, S. G. Brush, and C. W. F. Everitt, (Eds.). *Maxwell on heat and statistical mechanics* Bethlehem: Lehigh University Press, 1995.

- [Gearhart, 1990] C. A. Gearhart. Einstein before 1905: The early papers on statistical mechanics *American Journal of Physics*, 58, 468–480, 1990.
- [Gibbs, 1875] J. W. Gibbs. equilibrium of heterogenous substances. *Transactions of the Connecticut Academy*, 3, 103–246 and 343–524 (1878). Also in [Gibbs, 1906, pp. 55–353].
- [Gibbs, 1902] J. W. Gibbs. *Elementary principles in statistical mechanics*, New York, Scribner etc, 1902.
- [Gibbs, 1906] J. W. Gibbs. *The Scientific Papers of J. Willard Gibbs*, Vol. 1, Thermodynamics, Longmans, London, 1906. Reissued by Ox Bow Press, Woodbridge, Connecticut, 1993.
- [Giles, 1964] R. Giles. *Mathematical foundations of thermodynamics* Oxford: Pergamon, 1964.
- [Gold, 1956] T. Gold. Cosmic processes and the nature of time. In R. Colodny (Ed.), *Mind and Cosmos*. Pittsburgh: University of Pittsburgh Press, pp. 311–329, 1956.
- [Goldstein, 2001] S. Goldstein. Boltzmann’s approach to statistical mechanics. In J. Bricmont, D. Dürr, M.C. Galavotti, G. Ghirardi, F. Petruccione, and N. Zanghi (Eds.) *Chance in physics: foundations and perspectives*, Lecture Notes in Physics 574. Berlin: Springer-Verlag, pp. 39–54, 2001. Also as e-print cond-mat/0105242.
- [Gorini et al., 1976] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan. Completely positive dynamical semigroups of N -level systems. *Journal of Mathematical Physics* 17 8721–825, 1976.
- [Grimmett and Stirzaker, 1982] G. R. Grimmett and D. R. Stirzaker. *Probability and random processes*. Oxford: Clarendon Press, 1982.
- [de Groot, 1951] S. de Groot. *Thermodynamics of irreversible processes*. Amsterdam: North-Holland, 1951.
- [de Groot and Mazur, 1961] S. R. de Groot and P. Mazur. *Non-equilibrium thermodynamics*. Amsterdam: North-Holland, 1961.
- [Gross, 1997] D. H. E. Gross. Microcanonical thermodynamics and statistical fragmentation of dissipative systems. The topological structure of the N -body phase space. *Physics Reports*, 279, 119–201, 1997.
- [Gross and Votyakov, 2000] D. H. E. Gross and E. V. Votyakov. Phase transitions in “small” systems. *European Physical Journal B* 15, 115–126, 2000.
- [Grünbaum, 1973] A. Grünbaum. Is the coarse-grained entropy of classical statistical mechanics an anthropomorphism? In *Philosophical problems of space and time*, 2nd ed., Dordrecht: Reidel, pp. 646–665, 1973.
- [Guthrie, 1874] F. Guthrie. Molecular motion. *Nature*, 10 123, 19=874. Also in [Garber et al., 1995, p.143–145].
- [Greven et al., 2003] A. Greven, G. Keller, and G. Warnecke, (Eds.) *Entropy*. Princeton: Princeton University Press, 2003.
- [ter Haar, 1955] D. ter Haar. Foundations of statistical mechanics. *Reviews of Modern Physics*, 27, 289–338, 1955.
- [Hacking, 1975] I. Hacking. *The emergence of probability*. Cambridge: Cambridge University Press, 1975.
- [Hacking, 1990] I. Hacking. *The taming of Chance*. Cambridge: Cambridge University Press, 1990.
- [Van Harn and Holeyijn, 1991] K. van Harn and P. J. Holeyijn. *Markov-ketens in diskrete tijd*. Utrecht: Epsilon, 1991.
- [Herschel, 1850] J. F. W. Herschel. Quetelet on Probabilities *Edinburgh Review*. Also in *Essays from the Edinburgh and Quarterly Reviews with addresses and other pieces*, London: Longman, Brown, Green, Longmans and Roberts, 1857, pp. 365–465, 1850.
- [Hertz, 1910] P. Hertz. Über die mechanischen Grundlagen der Thermodynamik. *Annalen der Physik*, 33, 225–274; 537–552, 1910.
- [Hill, 1987] T. L. Hill. *Thermodynamics of small systems*. New York: Benjamin, 1987.
- [Höflechner, 1994] W. Höflechner. *Ludwig Boltzmann Leben und Briefe*. Graz: Akademische Druck- und Verlagsanstalt, 1994.
- [Huang, 1987] K. Huang. *Statistical mechanics*. New York: Wiley, 1987.
- [Huggett, 1999] N. Huggett. Atomic metaphysics. *Journal of Philosophy*, 96, 5–24, 1999.
- [Hutter and Wang, 2003] K. Hutter and Y. Wang. Phenomenological thermodynamics and entropy principles, 2002. In [Greven et al., 2003, pp. 55–78].
- [Illner and Neunzert, 1987] R. Illner and H. Neunzert. The concept of irreversibility in kinetic theory of gases. *Transport Theory and Statistical Physics*, 16, 89–112, 1987.

- [Janssen, 2002] M. Janssen. Dog fleas and tree trunks: the Ehrenfests marking the territory of Boltzmann's H -theorem. Unpublished, 2002.
- [Jauch, 1972] J. Jauch. On a new foundation of equilibrium thermodynamics. *Foundations of Physics*, 2, 327–332, 1972.
- [Jauch, 1975] J. Jauch. Analytical thermodynamics. Part 1. Thermostatistics—general theory. *Foundations of Physics*, 5, 111–132, 1975.
- [Jaynes, 1965] E. T. Jaynes. Gibbs vs. Boltzmann entropies. *American Journal of Physics*, 33, 391–398, 1965. Also in [Jaynes, 1983, pp. 77–86].
- [Jaynes, 1967] E. T. Jaynes. Foundations of probability theory and statistical mechanics. In: M. Bunge (Ed.) *Delaware seminar in the foundations of physics*. Berlin: Springer-Verlag, pp. 77–101, 1967. Also in [Jaynes, 1983, pp. 89–113].
- [Jaynes, 1983] E. T. Jaynes. *Probability, statistics and statistical physics*. R. Rosenkrantz (Ed.) Dordrecht: Reidel, 1983.
- [Jaynes, 1992] E. T. Jaynes. The Gibbs paradox. In C.R. Smith, G.J. Erickson, & P.O. Neudorfer, (Eds.) *Maximum entropy and Bayesian methods*. Dordrecht: Kluwer, pp. 1–22, 1992.
- [Jepps *et al.*, 2000] O. G. Jepps, G. Ayton, and D. J. Evans. Microscopic expressions for the thermodynamic temperature. *Physical Review E*, 62, 4757–4763, 2000.
- [Kadanoff, 2000] L. P. Kadanoff. *Statistical physics : statics, dynamics and renormalization* Singapore: World Scientific, 2000.
- [van Kampen, 1962] N. G. van Kampen. Fundamental problems in the statistical mechanics of irreversible processes. In E.G.D. Cohen (Ed.), *Fundamental problems in statistical mechanics* Amsterdam: North-Holland, pp.173–202, 1962.
- [van Kampen, 1981] N. G. van Kampen. *Stochastic processes in chemistry and physics*. Amsterdam: North-Holland, 1981.
- [van Kampen, 1984] N. G. van Kampen. The Gibbs paradox. In W.A. Parry (ed.), *Essays in theoretical physics in honour of Dirk ter Haar*, Oxford: Pergamon Press, pp. 303–312, 1984.
- [van Kampen, 1994] N. G. van Kampen. Models for dissipation in quantum mechanics. In J.J.Brey, J. Marro, J.M. Rubi, M. San Miguel (Eds.) *25 years of non-equilibrium statistical mechanics* Berlin: Springer, 1994.
- [van Kampen, 2002] N. G. van Kampen. The road from molecules to Onsager. *Journal of Statistical Physics*, 109, 471–481, 2002.
- [van Kampen, 2004] N. G. van Kampen. A new approach to noise in quantum mechanics *Journal of Statistical Physics*, 115, 1057–1072, 2004.
- [Karakostas, 1996] V. Karakostas. On the Brussels school's arrow of time in quantum theory. *Philosophy of Science*, 63, 374–400, 1996.
- [Kelly, 1979] F. P. Kelly. *Reversibility and stochastic networks* Chichester: Wiley, 1979. Also at http://www.statslab.cam.ac.uk/~afrb2/kelly_book.html.
- [Kestin, 1976] J. Kestin. *The second law of thermodynamics*. Stroudsburg, Pennsylvania: Dowden, Hutchinson and Ross, 1976.
- [Keynes, 1921] J. M. Keynes. *A treatise on probability*. London: Macmillan, 1921.
- [Khinchin, 1949] A. Khinchin. *Mathematical foundations of statistical mechanics*. (New York: Dover), 1949.
- [Kirchhoff, 1894] G. Kirchhoff. *Vorlesungen über mathematische Physik, Vol IV: Theorie der Wärme* M. Planck (Ed.) Leipzig: Teubner, 1894.
- [Kirsten and Körber, 1975] C. Kirsten and H.-G. Körber. *Physiker über Physiker*, Berlin: Akademie-Verlag, 1975.
- [Klein, 1970] M. J. Klein. Maxwell, his demon and the second law of thermodynamics, *American Scientist* 58, 82–95, 1970; also in [Leff and Rex, 1987, pp. 59–72].
- [Klein, 1972] M. J. Klein. Mechanical explanation at the end of the nineteenth century. *Centaureus*, 17, 58–82, 1972.
- [Klein, 1973] M. J. Klein. The Development of Boltzmann's Statistical Ideas. In E.G.D. Cohen & W. Thirring (eds.), *The Boltzmann equation*, Wien: Springer, pp. 53–106, 1973.
- [Klein, 1974] M. J. Klein. Boltzmann, monocycles and mechanical explanation. In R.J. Seeger & R.S. Cohen (Eds.), *Philosophical foundations of science; Boston Studies in the Philosophy of Science XI* Dordrecht: Reidel, pp. 155–175, 1974.
- [Klein, 1978] M. J. Klein. The thermostatistics of J. Willard Gibbs: atyranformation of thermodynamics. In E.G. Forbes (Ed.) *Human implications of scientific advance*. Edinburgh: Edinburgh University Press, pp. 314–330, 1978.

- [Kroes, 1985] P. A. Kroes. *Time: its structure and role in physical theories*. Dordrecht: Reidel, 1985.
- [Kurth, 1960] R. Kurth. *Axiomatics of classical statistical mechanics*. Oxford: Pergamon Press, 1960.
- [Ladyman et al., 2006] J. Ladyman, S. Presnell, T. Short, and B. Groisman. The connection between logical and thermodynamical irreversibility. <http://philsci-archive.pitt.edu/archive/00002374/>.
- [Landau and Lifshitz, 1987] L. E. Landau and E. M. Lifshitz. *Fluid Mechanics* 3rd Edition, Oxford: Butterworth-Heinemann, 1987.
- [Lanford, 1973] O. E. Lanford. Entropy and equilibrium states in classical statistical mechanics. In A. Lenard (Ed.) *Statistical mechanics and mathematical problems*. Berlin: Springer-Verlag, pp. 1–113.
- [Lanford, 1975] O. E. Lanford. Time evolution of large classical systems. In J. Moser (Ed.) *Dynamical Systems, Theory and Applications*, Lecture Notes in Theoretical Physics Vol. 38, Berlin: Springer, pp. 1–111, 1975.
- [Lanford, 1976] O. E. Lanford. On a derivation of the Boltzmann equation. *Astérisque*, 40 117–137, 1976. Also in J.L. Lebowitz & E.W. Montroll (Eds.) *Nonequilibrium phenomena I: the Boltzmann Equation*. Amsterdam: North-Holland, 1983.
- [Lanford, 1981] O. E. Lanford. The hard sphere gas in the Boltzmann-Grad limit. *Physica*, 106A, 70–76, 1981.
- [Lavis, 2004] D. A. Lavis. The spin-echo system reconsidered. *Foundations of Physics*, 34, 669–688, 2004.
- [Lavis, 2005] D. A. Lavis. Boltzmann and Gibbs: An attempted reconciliation. *Studies In History and Philosophy of Modern Physics*, 36, 245–273, 2005.
- [Lebowitz, 1994] J. Lebowitz. Time's arrow and Boltzmann's entropy. In J.J. Halliwell, J. Pérez-Mercader & W.H. Zurek (Eds.) *Physical origins of time asymmetry*. Cambridge: Cambridge University Press, pp. 131–146, 1994.
- [Lebowitz, 1999] J. L. Lebowitz. Statistical mechanics: A selective review of two central issues. *Reviews of Modern Physics*, 71, S346–S357, 1999. math-ph/0010018.
- [Lebowitz and Penrose, 1973] J. Lebowitz and O. Penrose. Modern ergodic theory. *Physics Today*, 26 23–29, 1973.
- [Leeds, 1989] S. Leeds. Discussion: D. Malament and S. Zabell on Gibbs phase averaging. *Philosophy of Science*, 56, 325–340, 1989.
- [Leff and Rex, 1987] H. S. Leff and A. F. Rex. *Maxwell's Demon 2*. Bristol: Institute of Physics, 1987.
- [Leinaas and Myrheim, 1977] J. M. Leinaas and J. Myrheim. On the theory of identical particles. *Il Nuovo Cimento*, 37B, 1–23, 1977.
- [Lieb, 1976] E. H. Lieb. The stability of matter. *Reviews of Modern Physics*, 48, 553–569, 1976.
- [Lieb and Lebowitz, 1973] E. H. Lieb and J. L. Lebowitz. Lectures on the thermodynamic limit for Coulomb systems. In A. Lenard (Ed.) *Statistical mechanics and mathematical problems*. Berlin: Springer-Verlag, pp. 136–162, 1973.
- [Lieb and Yngvason, 1999] E. Lieb and J. Yngvason. The physics and mathematics of the second law of thermodynamics, *Physics Reports*, 310, 1–96; erratum 314, (1999), p. 669, 1999. Also as e-print: <http://xxx.lanl.gov/abs/cond-mat/9708200>.
- [Lindblad, 1976] G. Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48, 119–130, 1976.
- [Lindblad, 1983] G. Lindblad. *Non-equilibrium entropy and irreversibility* Dordrecht: Reidel, 1983.
- [van Lith, 2001a] J. van Lith. Ergodic theory, interpretations of probability and the foundations of statistical mechanics *Studies in History and Philosophy of Modern Physics*, 32, 581–594, 2001.
- [van Lith, 2001b] J. van Lith. *Stirr in Stillness*. Ph.D. Thesis, Utrecht University, 2001.
- [Liu, 1999] C. Liu. Explaining the emergence of cooperative phenomena. *Philosophy of Science* 66 (*Proceedings*), S92–S106, 1999.
- [Lorentz, 1887] H. A. Lorentz. Über das Gleichgewicht der lebendige Kraft unter Gasmolekülen. *Wiener Berichte*, 95, 115–152, 187, 1887. Also in: *Collected Papers* The Hague: Martinus Nijhoff 1938, pp. 74–111.

- [Lorentz, 1916] H. A. Lorentz. *Les theories statistiques en thermodynamique*. Leipzig: Teubner, 1916.
- [Loschmidt, 1876] J. Loschmidt. Über die Zustand des Wärmegleichgewichtes eines Systems von Körpern mit Rücksicht auf die Schwerkraft. *Wiener Berichte*, 73, 128–142, 366–372 (1876), 75, 287–298, 76, 209–225 (1877).
- [Maes and Netočný, 2003] C. Maes and K. Netočný. Time-reversal and entropy. *Journal of Statistical Physics*, 110, 269–310, 2003.
- [Mackey, 1992] M. C. Mackey. *Time's arrow: the origins of thermodynamic behavior*. New York: Springer, 1992.
- [Mackey, 2001] M. C. Mackey. Microscopic dynamics and the second law of thermodynamics. In: *Time's Arrows, quantum Measurements and Superluminal Behavior*. C. Mugnai, A. Ranfagni & L.S. Schulman (Eds.) (Roma: Consiglio Nazionale delle Ricerche), 2001.
- [Malament and Zabell, 1980] D. B. Malament and S. L. Zabell. Why Gibbs Phase Averages Work — The Role of Ergodic Theory. *Philosophy of Science*, 56, 339–349, 1980.
- [Mandelbrot, 1956] B. Mandelbrot. An outline of a purely phenomenological theory of statistical thermodynamics. 1. canonical ensembles. *IRE Transactions on Information Theory*, IT-2, 190–203, 1956.
- [Mandelbrot, 1962] B. Mandelbrot. The role of sufficiency and of estimation in thermodynamics. *Annals of Mathematical Statistics*, 33, 1021–1038, 1962.
- [Mandelbrot, 1964] B. Mandelbrot. On the derivation of statistical thermodynamics from purely phenomenological principles. *Journal of Mathematical Physics*, 5, 164–171. 1964.
- [Mañé, 1987] R. Mañé. *Ergodic theory and differentiable dynamics*. Berlin: Springer, 1987.
- [Maroney, 2005] O. J. E. Maroney. The (absence of a) relationship between thermodynamic and logical reversibility. *Studies In History and Philosophy of Modern Physics*, 36, 355–374, 2005.
- [Martin-Löf, 1979] A. Martin-Löf. *Statistical mechanics and the foundations of thermodynamics*. Berlin: Springer-Verlag, 1979.
- [Maxwell, 1860] J. C. Maxwell. Illustrations of the dynamical theory of gases. *Philosophical Magazine*, 19, 19–32; 20, 21–37, 19=860. Also in [Garber *et al.*, 1986, pp. 285–318].
- [Maxwell, 1867] J. C. Maxwell. On the dynamical theory of gases. *Philosophical Transactions of the Royal Society of London*, 157 49–88 (1867). Also in [Brush, 1966] and [Garber *et al.*, 1986, pp. 419–472]
- [Maxwell, 1872] J. C. Maxwell. *Theory of heat*. (2nd ed.) London: Longmans, Green and Co, 1872.
- [Maxwell, 1873] J. C. Maxwell. Molecules. *Nature*, 8, 437–441 (1873). Also in [Garber *et al.*, 1986, pp. 138–154]
- [Maxwell, 1873] J. C. Maxwell. On the final state of a system of molecules in motion subject to forces of any kind. *Nature*, 8, 537–538 (1867), 1873. Also in [Garber *et al.*, 1995, pp. 138–143]
- [Maxwell, 1875] J. C. Maxwell. On the dynamical evidence of the molecular constitution of bodies. *Nature*, 11 357–359, 374–377, (1875). Also in [Garber *et al.*, 1986, pp. 216–237].
- [Maxwell, 1877] J. C. Maxwell. A treatise on the kinetic theory of gases by Henry William Watson, *Nature*, 18, 242–246, 1877. Also in [Garber *et al.*, 1995, p. 156–167].
- [Maxwell, 1878a] J. C. Maxwell. Constitution of bodies. *Encyclopaedia Britannica*, ninth edition, Vol. 6, pp. 310–313, 1878. Also in [Garber *et al.*, 1986, pp. 246–254].
- [Maxwell, 1878b] J. C. Maxwell. Diffusion. *Encyclopedia Britannica* ninth edition, Vol. 7, pp. 214–221, 1978; also in [Garber *et al.*, 1986, pp. 525–546].
- [Maxwell, 1878c] J. C. Maxwell. On stresses in rarified gases arising from inequalities of temperature *Proceedings of the Royal Society of London*, 27 304–308; *Philosophical Transactions of the Royal Society of London*, 170, (1880) 231–256, 1978. Also in [Garber *et al.*, 1995, pp. 357–386].
- [Maxwell, 1879] J. C. Maxwell. On Boltzmann's theorem on the average distribution of energy in a system of material points. *Transactions of the Cambridge Philosophical Society*, 12, 547–570, 1879. Also in [Garber *et al.*, 1995, pp. 357–386].
- [Mehra, 1998] J. Mehra. Josiah Willard Gibbs and the Foundations of Statistical Mechanics *Foundations of Physics*, 28, 1785–1815, 1998.
- [Mehra and Sudarshan, 1972] J. Mehra and E. C. G. Sudarshan. Some reflections on the nature of entropy, irreversibility and the second law of thermodynamics, *Nuovo Cimento B*, 11, 215–256, 1972.
- [Meixner, 1969] J. Meixner. Processes in simple thermodynamic materials *Archive for Rational Mechanics and Analysis* 33, 33–53, 1969.

- [Meixner, 1970] J. Meixner. On the foundations of thermodynamics of processes. In B. Gal-Or, E.B. Stuart & A. Brainard (Eds.), *A Critical Review of Thermodynamics*. Baltimore: Mono Book Corporation, pp. 37–47, 1970. Also in [Kestin, 1976], pp. 313–323.
- [Moran, 1961] P. A. P. Moran. Entropy, Markov processes and Boltzmann's H-theorem. *Proceedings of the Cambridge Philosophical Society*, 57, 833–842, 1961.
- [Morrison, 2000] M. Morrison. *Unifying scientific theories: physical concepts and mathematical structures*. Cambridge: Cambridge University Press. 2000.
- [Morrison, 1966] P. Morrison. Time's arrow and external perturbations. In A. de Shalit, H. Feshbach, & L. van Hove (Eds.), *Preludes in Theoretical Physics in honor of V. F. Weisskopf*. Amsterdam: North Holland, pp. 347–, 1966.
- [Müller, 2003] I. Müller. Entropy in non-equilibrium. In [Greven et al., 2003, pp. 79–107].
- [Navarro, 1998] L. Navarro. Gibbs, Einstein and the foundations of statistical mechanics. *Archive for History of Exact Science*, 53, 147–180, 1998.
- [Nemytskii and Stepanov, 1960] Nemytskii, V.V., & V. V. Stepanov. *Qualitative theory of differential equations*. Princeton: Princeton University Press, 1960.
- [Norton, 2005] J. D. Norton. Eaters of the lotus: Landauer's principle and the return of Maxwell's demon. *Studies In History and Philosophy of Modern Physics*, 36, 375–411, 2005.
- [von Neumann, 1932] J. von Neumann. Proof of the quasi-ergodic hypothesis. *Proceedings of the National Academy of sciences of the United States of America*, 18, 70–82 and 263–266, 1932.
- [Obcemea and Brändas, 1983] Ch. Obcemea and E. Brändas. Analysis of Prigogine's theory of subdynamics. *Annals of Physics*, 147, 383–430, 1983.
- [Olsen, 1993] E. T. Olsen. Classical mechanics and entropy. *Foundations of Physics Letters*, 6, 327–337, 1993.
- [Penrose, 1970] O. Penrose. *Foundations of statistical mechanics: a deductive treatment*. Oxford: Pergamon Press, 1970.
- [Penrose, 1979] O. Penrose. Foundations of statistical mechanics. *Reports on Progress in Physics*, 42, 1937–2006, 1979.
- [Penrose and Percival, 1962] O. Penrose and I. Percival. The direction of time. *Proceedings of the Physical Society*, 79, 605–616, 1962.
- [Pešić, 1991] P. D. Pešić. The principle of identity and the foundations of quantum theory. I. The Gibbs paradox. *American Journal of Physics*, 59, 971–974, 1991.
- [Petersen, 1983] K. Petersen. *Ergodic theory*. Cambridge: Cambridge University Press, 1983.
- [Pitowsky, 2001] I. Pitowsky. Local fluctuations and local observers in equilibrium statistical mechanics. *Studies In History and Philosophy of Modern Physics*, 32, 595–607, 2001.
- [Plancherel, 1913] M. Plancherel. Foundations of statistical mechanics. *Annalen der Physik*, 42, 1061–1063, 1913.
- [von Plato, 1991] J. von Plato. Boltzmann's ergodic hypothesis. *Archive for History of Exact Sciences*, 42, 71–89, 1991.
- [von Plato, 1994] J. von Plato. *Creating modern probability*. Cambridge: Cambridge University Press, 1994.
- [Poincaré, 1889] H. Poincaré. Sur les tentatives d'explication mécanique des principes de la thermodynamique. *Comptes Rendus de l'Académie des Sciences (Paris)*, 108, 550–553, 1889. English translation in [Olsen, 1993].
- [Poincaré, 1893] H. Poincaré. Le mécanisme et l'expérience. *Revue de Métaphysique et de Morale*, 1, 534–537, 1893. English translation in [Brush, 1966].
- [Poincaré, 1896] H. Poincaré. *Calcul des Probabilités*. Paris: Carré, 1896.
- [Popper, 1982] K. Popper. *Quantum theory and the schism in physics*. London: Hutschinson, 1982.
- [Price, 1996] H. Price. *Time's arrow and Archimedes' point*. New York: Oxford University Press, 1996.
- [Prigogine, 1955] I. Prigogine. *Introduction to the thermodynamics of irreversible processes*. New York: Interscience, 1955.
- [Ray, 1984] J. R. Ray. Correct Boltzmann counting. *European Journal of Physics*, 5 219–224, 1984.
- [Redhead, 1995] M. Redhead. *From physics to metaphysics*. Cambridge: Cambridge University Press, 1995.
- [Reichenbach, 1956] H. Reichenbach. *The direction of Time*. Berkeley: University of California Press, 1956.

- [Ridderbos, 2002] T. M. Ridderbos. The coarse-graining approach to statistical mechanics: how blissful is our ignorance? *Studies in History and Philosophy of Modern Physics*, 33, 65–77, 2002.
- [Ridderbos and Redhead, 1998] T. M. Ridderbos and M. L. G. Redhead. The spin-echo experiment and the second law of thermodynamics. *Foundations of Physics*, 28, 1237–1270, 1998.
- [Rosenthal, 1913] A. Rosenthal. Beweis der Unmöglichkeit ergodischer mechanischer Systeme. *Annalen der Physik*, 42, 796–806, 1913. English translation in [Brush, 2003, p.505–523].
- [Rovelli, 2006] C. Rovelli. This volume, chapter 12, 2006.
- [Ruelle, 1969] D. Ruelle. *Statistical mechanics: rigorous results*. New York: Benjamin, 1969.
- [Rugh, 1997] H. H. Rugh. Dynamical approach to temperature *Physical Review Letters* 78, 772–774, 1997.
- [Rugh, 2001] H. H. Rugh. Microthermodynamic formalism *Physical Review E* 64, 055101, 2001.
- [Saunders, 2006] S. Saunders. On the explanation for quantum statistics. *Studies in History and Philosophy of Modern Physics*, 37, 2006.
- [Schrödinger, 1950] E. Schrödinger. Irreversibility. *Proceedings of the Royal Irish Academy* 53, 189–195, 1950.
- [Simányi and Szász, 1999] N. Simányi and D. Szász. Hard balls systems are completely hyperbolic. *Annals of Mathematics*, 149, 35–96, 1999.
- [Sinai, 1963] Ya. G. Sinai. On the foundation of the ergodic hypothesis for a dynamical system of statistical mechanics *Soviet Mathematics Doklady*, 4, 1818–1822, 1963.
- [Sinai and Chernov, 1987] Ya. G. Sinai and N. I. Chernov. Ergodic properties of certain systems of $2 - D$ discs and $3 - D$ balls. *Russian Mathematical Surveys*, 42 181–207. Also in Ya.G. Sinai (Ed.) *Dynamical systems; collection of papers* Singapore: World Scientific (1991) pp. 345–371, 1987.
- [Shenker, 2000] O. Shenker. Interventionism in statistical mechanics: some philosophical remarks. <http://philsci-archive.pitt.edu/archive/00000151/> 2000.
- [Sklar, 1973] L. Sklar. Statistical explanation and ergodic theory. *Philosophy of Science*, 40, 194–212, 1973.
- [Sklar, 1993] L. Sklar. *Physics and Chance. Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge: Cambridge University Press, 1993.
- [Sklar, 2002] L. Sklar. *Theory and truth*. Oxford: Oxford University Press, 2002.
- [Spohn, 1980] H. Spohn. Kinetic equations from Hamiltonian dynamics: Markovian limits. *Reviews of Modern Physics* 52, 569–615, 1980.
- [Spohn, 1991] H. Spohn. *Large Scale Dynamics of interacting Particles*. Berlin: Springer, 1991.
- [Streater, 1995] R. F. Streater. *Statistical dynamics; a stochastic approach to non-equilibrium thermodynamics*. London: Imperial College Press, 1995.
- [Styer, 2004] D. F. Styer. What good is the thermodynamic limit?" *American Journal of Physics* 72, 25–29; Erratum p. 1110, 2004.
- [Sudarshan et al., 1961] E. C. G. Sudarshan, P. M. Mathews, and J. Rau. Stochastic dynamics of quantum-mechanical systems, *Physical Review*, 121, 920–924, 1961.
- [Szász, 1996] D. Szász. Boltzmann's ergodic hypothesis, a conjecture for centuries? *Studia Scientiarum Mathematicarum Hungaria*, 31, 299–322, 1996.
- [Szilard, 1925] L. Szilard. Über die Ausdehnung der phänomenologischen Thermodynamik auf die Schwankungserscheinungen. *Zeitschrift für Physik*, 32, 753–788, 1925.
- [Tabor, 1989] M. Tabor. *Chaos and integrability in nonlinear dynamics: an introduction*. New York: Wiley, 1989.
- [Thomsen and Hartka, 1962] J. S. Thomsen and Th. J. Hartka. Strange Carnot cycles; thermodynamics of a system with a density extremum *American Journal of Physics* 30, 26–33, 1962.
- [Thompson, 1972] C. Thompson. *Mathematical Statistical Mechanics*. Princeton: Princeton University Press, 1972.
- [Tisza, 1966] L. Tisza. *Generalized Thermodynamics*. Cambridge, Mass.: MIT Press, 1966.
- [Tisza and Quay, 1963] L. Tisza and P. M. Quay. The statistical thermodynamics of equilibrium. *Annals of Physics*, 25, 48–90, 1963. Also in [Tisza, 1966, pp. 245–287].
- [Tolman, 1938] R. C. Tolman. *The principles of statistical mechanics*. London: Oxford University Press, 1938.
- [Touchette, 2003] H. Touchette. *Equivalence and nonequivalence of the microcanonical and canonical ensembles: a large deviations study*. Ph.D Thesis, McGill University, Montréal, 2003.

- [Touchette et al., 2004] H. Touchette, R. S. Ellis, and B. Turkington. An introduction to the thermodynamic and macrostate levels of nonequivalent ensembles. *Physica A* 340, 138-146, 2004.
- [Truesdell, 1961] C. Truesdell. Ergodic theory in classical statistical mechanics. In P. Caldirola (ed.), *Ergodic theories*, New York: Academic Press, 1961.
- [Truesdell, 1969] C. Truesdell. *Rational thermodynamics*. New York: McGraw-Hill, 1969.
- [Truesdell, 1980] C. Truesdell. *The tragicomical history of thermodynamics 1822–1854*. New York: Springer, 1980.
- [Uffink, 1995] J. Uffink. Can the maximum entropy principle be regarded as a consistency requirement? *Studies in History and Philosophy of Modern Physics*, 26, 223–261, 1995.
- [Uffink, 1996] J. Uffink. The constraint rule of the maximum entropy principle. *Studies in History and Philosophy of Modern Physics*, 27, 47–79, 1996.
- [Uffink, 1996b] J. Uffink. Nought but molecules in motion. *Studies In History and Philosophy of Modern Physics*, 27, 373–387, 1996.
- [Uffink, 2001] J. Uffink. Bluff your way in the second law of thermodynamics. *Studies in History and Philosophy of Modern Physics*, 32, 305–394, 2001.
- [Uffink, 2003] J. Uffink. Irreversibility and the second law of thermodynamics. In [Greven et al., 2003, pp. 121–146], 2003.
- [Uffink, 2004] J. Uffink. Boltzmann's Work in Statistical Physics. In E.N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2004 Edition), <http://plato.stanford.edu/archives/win2004/entries/statphys-Boltzmann>.
- [Uffink, 2005] J. Uffink. Rereading Ludwig Boltzmann. In P. Hájek, L. Valdés-Villanueva & D. Westerståhl (Eds.) *Logic, Methodology and Philosophy of Science. Proceedings of the twelfth international congress*. London: King's College Publications, pp. 537–555, 2005.
- [Uffink, 2006] J. Uffink. Insuperable difficulties: Einstein's statistical road to molecular physics. *Studies in History and Philosophy of Modern Physics*, 2006.
- [Uffink and van Lith, 1999] J. Uffink and J. van Lith. Thermodynamic uncertainty relations. *Foundations of Physics*, 29, 655–692, 1999.
- [Uhlenbeck and Ford, 1963] G. E. Uhlenbeck. and G. W. Ford. *Lectures in statistical mechanics*. Providence, Rhode Island: American Mathematical Society, 1963.
- [Vranas, 1998] P. B. M. Vranas. Epsilon-ergodicity and the success of equilibrium statistical mechanics. *Philosophy of Science*, 65, 688–708, 1998.
- [Wald, 2001] R. M. Wald. The thermodynamics of black holes. *Living Reviews in Relativity*, 4, 6. <http://www.livingreviews.org/lrr-2001-6>.
- [Winsberg, 2004] E. Winsberg. Laws and statistical mechanics. *Philosophy of Science*, 71, 707–718, 2004.
- [Yourgrau et al.,] abc W. Yourgrau, A. van der Merwe, and G. Raw. *Treatise on irreversible thermophysics*, New York: Macmillan, 1966.
- [Zermelo, 1896a] E. Zermelo. Ueber einen Satz der Dynamik und die mechanische Wärmetheorie. *Annalen der Physik*, 57, 485–494, 1896. English translation in [Brush, 2003, pp. 382–391].
- [Zermelo, 1896b] E. Zermelo. Ueber mechanische Erklärungen irreversibler Vorgänge. *Annalen der Physik*, 59, 793–801, 1896. English translation in [Brush, 2003, pp. 403–411].

QUANTUM STATISTICAL PHYSICS

G erard G. Emch

1 INTRODUCTION

In search for the headwaters of the Missouri during their 1804–1806 expedition, Lewis and Clark decreed that the river begins at the confluence of three streams — the Jefferson, Gallatin, and Madison rivers — and ends as a main tributary to the mighty Mississippi.

Similarly, and with some of the same arbitrariness, three major headings can be used to mark the beginnings of quantum statistical physics (QSP): Planck’s “quantum hypothesis” following his 1900 papers [Planck, 1900a; Planck, 1900b], Gibbs’ 1902 book on “statistical mechanics” [Gibbs, 1902], and what is now known as Einstein’s 1905 “Brownian motion” [Einstein, 1905b]. Pushing the metaphor into our own days, the power of QSP is manifest in the landscape of condensed matter physics (from solid state physics to astrophysics). The navigation there, albeit often tentative, has brought to shore predictions that have been confirmed with a precision impressive enough to clamor for a consistent explanation. The purpose of this chapter is to point to directions along which such explanations may be found. I begin this search by tracing briefly the course of the three tributaries mentioned above, thus by identifying the initial motivations for QSP.

Planck’s long-lasting hesitations indicate how much in advance he was, not only of his own time, but perhaps even of himself; e.g., at first, he had put forward his black-body radiation law on account of the nature of the *body* — little oscillators in the walls — rather than on account of the nature of the *radiation*. As Planck was transposing to the description of electromagnetic waves the counting arguments Boltzmann used in the thermal physics of material bodies, he initially left open the question of whether this was a mere formal analogy, or whether it was one that could be justified from putative interactions between radiation and matter; or whether, yet, this speculative analogy had deeper roots. Planck’s reluctance still shows through in the recommendation he wrote in 1913 to support young Einstein’s early election to the Prussian Academy of Sciences: “That he may sometimes have missed the target of his speculations, as for example in his hypothesis of the light quanta, cannot really be held against him.” While this may be seen as a barb directed to [Einstein, 1905a], note nevertheless that Planck’s reference to a quantum *hypothesis* is not a passing accident: he was meticulous in his use of words; consider, for instance the use of “theory”, “theorem”, and “hypothesis”

in the title of his 1911 address to the German Chemical Society [Planck, 1911]. Soon thereafter, the rest of the world overcame his scruples: the Nobel prize was awarded to Planck in 1918 for “his discovery of energy quanta”; and to Einstein in 1921 for “his discovery of the law of the photoelectric effect.” For each of them, the *laudatio* calls attention to their respective contributions to the nascent QSP, specifically: the black body radiation for Planck and the specific heat of solids for Einstein; see subsections 2.1 and 2.3 below.

Gibbs’ book [Gibbs, 1902] focuses on *classical* statistical physics. While the basic concepts had been apprehended differently by the German Clausius, the Austrian Boltzmann and the British Maxwell, the American Gibbs proposes that the field has reached the modicum of maturity necessary for a consolidation of the foundations; for axiomatization in other fields, compare with Hilbert [Hilbert, 1900; Hilbert, 1899; Hilbert, 1918], and Einstein [Einstein, 1921]. Even in the classical context, Gibbs’ reluctance to invoke Boltzmann’s ergodic postulate points to the persistence of unresolved issues regarding what Gibbs calls in the very title of his book “the rational foundation of thermodynamics”; for a brief presentation of those aspects of Gibbs’ work that may be most relevant to my purpose here, see [Uffink, 2006, section 5]. It pertains to the present chapter to examine how much of this dichotomy persists in the quantum realm, and the extent to which whatever persists is relevant to the explanatory purposes of QSP.

Einstein’s papers on Brownian motion still reside conceptually in the realm of classical physics. In spite of the neglect in which many mathematicians still held the foundations of probabilistic theories around the turn of the twentieth century (cf. e.g. [Hilbert, 1900, Problem 6]), Einstein’s approach stands as a witness to the fact that stochastic arguments — i.e. arguments involving random processes — had gained currency in the physicists’ marketplace. Einstein’s conclusions were widely (if not universally) accepted at face value as empirical proof of the existence of molecules, as not just computationally convenient small entities or units, but as objects with definite dimensions [Einstein, 1906b]. Furthermore, Einstein’s papers were not the isolated manifestation of a singular genius that the cumbersome title of his first paper might suggest [Einstein, 1905b]. On the one hand, from the physicist’s perspective, it must be noted that Einstein begins his second paper in the sequence with an acknowledgment that he had ignored the earlier contributions of Siedentopf and Gouy who had interpreted the “so-called Brownian motion” [Einstein *dixit*] as caused by the irregular thermal motions of the molecules [Einstein, 1906c; Gouy, 1888]. On the other hand, the modern mathematician will recognize, with the hindsight of practitioners such as Kac and Chandrasekhar, that Smolukowski simultaneously distilled from the same empirical sources the mathematical intuition allowing him to post a claim on what was to become the theory of stochastic processes [Smolukowski, 1906a; Smolukowski, 1916]. Yet, it was only in 1933 that Kolmogorov made precise the essentials of the underlying syntax, namely the mathematical theory of probability [Kolmogorov, 1933]. Even so, an unresolved issue remains to this day as to the proper semantics: von Mises’ collectives [von Mises, 1928] or de Finetti’s subjec-

tive assignments [de Finetti, 1937]. I post my stakes — see subsection 3.1 — on the latter issue when considering the extension of the theory of probability to the quantum realm, with special regard to the specific demands of QSP.

As this essay opens, the question arises as to whether the confluence of three streams of interest compounds the foundational problems of each of them or, on the contrary, whether they can be brought to inform one another. I aim my argument towards the latter view, although I am not oblivious to such ubiquitous problems as questioning what elements of reality should — or should not — be ascribed to individual microscopic quantum systems. As part of the larger problem of the reduction of thermodynamics by statistical mechanics, I consider specifically the question whether and how QSP can claim to explain the collective properties of many-body systems: it does postulate a quantum description at the microscopic level, while it has not obtained as yet an ontological grasp of the individual components of these systems. In my presentation I follow Einstein's admonition: "If you want to find out anything from the theoretical physicists about the methods they use ... don't listen to their words, fix your attention on their deeds." [Einstein, 1933].

2 EARLY SUCCESSES

In [Jammer, 1966] Max Jammer provides much of the specific historical documentation pertaining to the beginnings of quantum theory; and he discusses some of the ensuing debates in [Jammer, 1974]. Here, I start with a discussion of the early pragmatic successes of QSP, with special attention to two aspects: their classical mooring in the high temperature regime; and the understanding QSP gives of the particle-wave duality. Both of these aspects illustrate the added insight gained from the contextual differences coloring the answers to the same questions when asked in QSP rather than in the quantum theories of, say, the Bohr atom or scattering processes; compare with Mara Beller's perspective on the making of the quantum revolution [Beller, 1999].

2.1 Planck's interpolating formula for black-body radiation

The *experimental evidence* available to Planck was the spectral density $\rho_T(\nu)$ of the energy per unit volume of electromagnetic radiation, as a function of its frequency ν , when electromagnetic radiation is in equilibrium with a black-body at temperature T . In [Planck, 1900a; Planck, 1900b], Planck proposes to fit these data with the formula

$$(1) \quad \rho_T(\nu) = A \frac{h\nu}{e^{\frac{h\nu}{kT}} - 1} \quad \text{with} \quad A = \frac{8\pi\nu^2}{c^3}$$

where c is the speed of light, $k = R/N_{Av}$ is the Boltzmann constant, R is the universal gas constant and N_{Av} is the Avogadro number. In addition, a *new*

constant enters the formula, h , known nowadays as the Planck constant. While Planck himself would pretend that (1) had been a “lucky guess” such a formula could not have come into existence in a conceptual vacuum.

Two qualitative laws had been identified by Wien [1894], Stefan [1879] and Boltzmann [1884]. The *Wien displacement law* states that

$$(2) \quad \rho_T(\nu) = \nu^3 f\left(\frac{\nu}{T}\right)$$

where f is some undetermined function, satisfying the condition that the following integral converges

$$(3) \quad \frac{1}{V}E(T) = \int_0^\infty d\nu \rho_T(\nu)$$

which expresses the density, per unit volume, of the energy of the radiation at temperature T . Upon inserting (2) in (3), one receives immediately the *Stephan–Boltzmann law*:

$$(4) \quad E(T) = \sigma T^4$$

where σ is a constant. Planck’s proposal complies with these laws.

Two analytic expressions (or “laws”) had been proposed, which specify the function f in (2). One law, due to Wien [1896], reads:

$$(5) \quad \rho_T(\nu) = \alpha \nu^3 \exp^{-\gamma \frac{\nu}{T}} \quad .$$

With α and γ being two constants, this law had been confirmed empirically in the range where ν/T is *large*. In contrast, the other law, due to Rayleigh [1900], see also Jeans [1905a], is:

$$(6) \quad \rho_T(\nu) = \frac{8\pi}{c^3} \nu^2 kT$$

which had been confirmed empirically in the range where ν/T is *small*.

Clearly, (1) interpolates analytically between the Wien and Rayleigh–Jeans formulas; and it gives a quantitative meaning to the conditions that ν/T be “large” (resp. “small”), namely $\nu/T \gg k/h$ (resp. $\nu/T \ll k/h$). In the intermediate range, Planck’s interpolating formula fits experimental results very well, both qualitatively and quantitatively.

Planck’s colleagues could not fail to be impressed and Planck’s triumph would have been complete had he been able to explain his formula from first principles, at least to the considerable extent with which (2) to (6) could be understood. Instead, Planck has to resort to “an act of desperation” — his own words [Jammer, 1966] — and he constructs, after several attempts, a heuristic model in which the radiation exchanges energy in discrete *quanta* with putative “resonators” in thermodynamical equilibrium within the walls. The model suffers from several shortcomings — among them Planck’s adaptation of Boltzmann’s counting — and much uncertainty concerning its theoretical status:

Walter Nernst ... initially disliked quantum theory, claiming that it was ‘really nothing else than an interpolation formula ... only a rule for calculations ... but has proven so fruitful by the work of Planck ... and ... of Einstein that the duty of science is to take it seriously and to subject it to careful investigations’. [Jammer, 1966, p. 59]

The consensus that later settled in the physics community is that *any* attempt — Planck’s included — to derive (1) from first principles would be doomed to failure: (1) *is* a fundamental or primary law, i.e. one that is not to be explained, but the consequences of which ought to be explored.

2.2 Einstein’s fluctuation formula and the particle-wave duality

For a start, Einstein notes two shortcomings in Planck’s derivation. The first is formal, but nevertheless essential: Planck’s account does not conform to Boltzmann’s statistical counting as closely as Planck suggests. The second is pointed out in [Einstein, 1906a]: Planck’s treatment involves an inconsistency between: (a) his use of the (classical) Maxwell theory of electromagnetism to compute the average energy of a resonator in a radiation field; and (b) the assumption that the energy of a resonator can change only discontinuously. Together with other empirical problems — among which the photoelectric effect [Einstein, 1905a] — these difficulties led Einstein to propose that, while Planck’s radiation formula (1) has incontestable empirical merits, the “quantization” itself is to be looked for in the radiation field rather than in a dubious mechanism of interaction with the walls. Einstein’s criticism raises, in the same volley, the question of whether light is wave-like as accounted for by Maxwell’s electromagnetic theory; or whether it is particle-like as Newton’s theory had it before its purported falsification in interference experiments conducted in the early nineteenth century .

Einstein’s fluctuation formula [Einstein, 1909a] proposes that light should be viewed *simultaneously* as *both* particle *and* wave; specifically:

SCHOLIUM 1. Let the Planck’s spectral density $\rho_T(\nu)$ in (1) be interpreted as the average energy $\langle u_T(\nu) \rangle$ of quantum oscillators of frequency ν of the radiation in thermal equilibrium at temperature T . Then for all values of $h\nu/kT$, the energy fluctuation $\langle (\Delta u)^2 \rangle = kT^2 \partial_T \langle u_T(\nu) \rangle$ is the sum of two terms

$$\langle (\Delta u)^2 \rangle = \langle (\Delta u)^2 \rangle_p + \langle (\Delta u)^2 \rangle_w \quad \text{where}$$

$$(7) \quad \left\{ \begin{array}{l} \langle (\Delta u)^2 \rangle_p = \langle u_T(\nu) \rangle h\nu \\ \langle (\Delta u)^2 \rangle_w = \langle u_T(\nu) \rangle^2 \frac{c^3}{8\pi\nu^2} \end{array} \right\} \quad \text{and} \quad \langle (\Delta u)^2 \rangle_p / \langle (\Delta u)^2 \rangle_w = \exp \frac{h\nu}{kT} - 1$$

Hence, the particle-like contribution $\langle (\Delta u)^2 \rangle_p$ dominates when $h\nu/kT \gg 1$, and the wave-like contribution $\langle (\Delta u)^2 \rangle_w$ dominates when $h\nu/kT \ll 1$. In this in-

terpretation, *the particle-wave duality is thus a matter of degree*, rather than an alternative between the two mutually exclusive horns of a dilemma.

Less of a conceptual problem in QSP, this duality becomes more difficult to master in other empirical contexts where one may prefer to view a photon *either* as a particle *or* as a wave packet. Moreover, this duality has since been extended to all (sub-atomic) particles; e.g. phenomena usually associated with waves, such as diffraction of beams of light, have been observed as well with beams of electrons and then neutrons; cf. e.g. [Jammer, 1966, pp. 249-253]; or for an update [Rauch, 2005]. In other circumstances, one prefers to use a particle language, as for instance in the description of the photo-electric effect [Einstein, 1905a]; as reported in most QM textbooks, a photon impinging on a metallic surface causes an electron to be expelled; or in atomic spectroscopy, a particle — the atom — emits a beam of light; cloud- and bubble-chambers have since let us visualize interparticle collisions; and yet their description in scattering theory uses the so-called wave operator; cf. e.g. [Amrein *et al.*, 1977]. In the light of this duality, and following upon the speculations of Einstein and de Broglie, physicists have learned to adapt their language to the aspect they wish to emphasize. Yet, the persistent arguments about “self-interference” show that some residual ambiguities have yet to be resolved; cf. the long debate extending from [Taylor, 1909] to [Aichele *et al.*, 2005], and surely beyond.

Upon returning to the early manifestations of QSP, one ought to mention that the Einstein fluctuation formula (7) above, as well as the explanation of the temperature dependence of the specific heat of solids — see subsection 2.3 below — motivate the Ehrenfests’ suggestion [Ehrenfest and Ehrenfest, 1911] that in statistical mechanics, quantum behaviour manifests itself mostly at low temperatures, whereas classical behaviour emerges at high temperatures. The fact is that in many expressions, such as the Planck distribution (1), the Planck constant h and the temperature T appear together in a factor h/T , or in the form used in the sequel, $\hbar\beta$; hence in these expressions the “classical limit” $h \rightarrow 0$ and the “high-temperature limit” $T \rightarrow \infty$ are included in $(\hbar\beta) \rightarrow 0$. All refer to cases where the relevant energies, or energy densities, are extremely large when measured in the scale determined by the numerical value of the Planck constant.

2.3 *Debye’s specific heat of solids below the classical regime*

For the purpose of this subsection, the situation down in the field is that Dulong & Petit (1819) had proposed an argument to the effect that the specific heat — measured in calories per mole per degree — ought to be the same for all solids: $3R$ where R is the universal gas constant. Yet, it later became apparent that this “constant” could decrease dramatically with temperature, so much so that by the end of the nineteenth century, the experimental data led to the *conjecture* that the specific heat of solids becomes vanishingly small as the temperature approaches absolute $0K$. In the meantime, the discovery of X-rays by Roentgen (1895) had allowed several experimentalists — Ewald (1911), and at the suggestion of von

Laue, Friedrich and Knipping (1912) — to obtain diffraction patterns corroborating speculations that crystalline solids are regular lattices, at the vertices of which sit the atoms.

As no classical explanation of the observed drastic temperature dependence of the specific heat seemed forthcoming, Einstein and Debye offered the following model; cf. [Einstein, 1907; Einstein, 1911b; Debye, 1912].

The starting point is (1) above, the Planck formula for black-body radiation, now reinterpreted in terms of the vibrational modes of a solid at temperature T :

$$(8) \quad U(T) = \int d\nu g(\nu) U(\nu, T) \quad \text{with} \quad U(\nu, T) = \frac{h\nu}{e^{\frac{h\nu}{kT}} - 1} \quad \text{and} \quad \int_0^\infty d\nu g(\nu) = 3N$$

where N is the number of 3-dimensional oscillators in the solid. Where Einstein had assumed that g is concentrated on a fixed frequency ν_o , Debye chooses for g the simplest vibrational distribution that takes into account that in a crystal, say of volume V , the vibrations have a minimal wavelength of the order of the interatomic distance in the lattice:

$$(9) \quad g(\nu) = G \left\{ \begin{array}{ll} 1 & \text{if } 0 \leq \nu \leq \nu_o \\ 0 & \text{if } \nu > \nu_o \end{array} \right\} \quad \text{with} \quad G = \frac{12\pi\nu^2}{s^3} V \quad .$$

G takes into account that vibrations are now sound waves rather than electromagnetic waves — compare with $A = \frac{8\pi\nu^2}{c^3}$ in (1) — thus s is now the speed of sound, instead of the speed c of light; and the replacement of $8\pi = 2 \cdot 4\pi$ by $12\pi = (2+1) \cdot 4\pi$ reflects the fact that sound-waves in solids have, in addition to the two transverse polarizations also present in light, a third degree of freedom, namely longitudinal modes. These hypotheses entail the following consequence.

SCHOLIUM 2. There exists a temperature Θ , such that the specific heat satisfies

$$(10) \quad C_V \simeq \begin{cases} 3R & \text{for } T \gg \Theta \\ \frac{12}{5}\pi^4 R \left(\frac{T}{\Theta}\right)^3 & \text{for } T \ll \Theta \end{cases} .$$

Hence, the Debye model differentiates between two regimes: at high temperatures it recovers the Dulong–Petit law; and it predicts that as the temperature approaches 0 K , the specific heat vanishes according to $C_V \sim T^3$. In this model, the temperature Θ , now called the *Debye temperature*, depends on the solid considered through the cut-off frequency ν_o , and thus on the speed of sound in that solid and on its density N/V . The numerical value of Θ gives a quantitative estimate — for details, see subsection 6.1 — of what is meant by high and low temperature regimes for the specific heat of crystalline solids. Moreover, in Debye's model, C_V decreases monotonically and continuously over the whole range of temperatures $T \in \mathbb{R}^+$.

As a last comment on the passage from (1) to (8), note that by analogy with the *photons* as the quanta of light, the elementary sound vibrations in solids are viewed as quanta, now known as *phonons*.

2.4 BE-condensation: the long haul

When taking seriously the idea that the microscopic picture of the macroscopic world may be a quantum one, the most immediate question is to obtain the corresponding description of a quantum ideal gas; this came to be known as the Bose—Einstein gas, or simply the Bose gas [Bose, 1924; Einstein, 1924]. The starting point is the grand canonical partition function $Z(\Lambda, T, \mu)$ of an assembly of identical massive particles of mass m in equilibrium at temperature T and chemical potential μ ; this assembly is enclosed in a cubical box of volume $\Lambda \subset \mathbb{R}^3$, with periodic boundary conditions. As these particles are non-interacting, the total energy is the sum of their individual energies $\epsilon_k = \hbar^2 |k|^2 / 2m$, where $k \in Z^3$. The quantum hypothesis is that the Planck distribution (1) applies here so as to entail (with $\beta = 1/kT$) :

$$(11) \quad Z(\Lambda, T, \mu) = \prod_{k \in Z^3} (1 - \exp^{-\beta(\epsilon_k - \mu)})^{-1} .$$

From this formula, one computes the specific volume v and the pressure P , according to the rules learned in classical statistical mechanics; the so-called activity is defined as $z = \exp(\beta\mu)$:

$$(12) \quad v^{-1} = z \partial_z \frac{1}{|\Lambda|} \ln Z(\Lambda, T, \mu) \quad \text{and} \quad \beta P = \frac{1}{|\Lambda|} \ln Z(\Lambda, T, \mu) .$$

The problem is thus stated completely, although the consequences of (11–2.12) are not easy objects to compute directly. The solution involves a mathematical excursion through some classical analysis, and the reward is a nice physical bounty: a phase transition with the onset of a condensed phase at very low temperatures; not your classical ideal gas!

The necessary classical analysis — now widely available, cf. e.g. [Whittaker and Watson, 1927, p.280, ex. 7], [Erd elyi, 1953, I, pp. 27–30], or for some historical perspective [Truesdell, 1945] — was already known to our pioneers, and they did recognize that in the limit $\Lambda \uparrow \mathbb{R}^3$, these sums reduce to:

$$(13) \quad \left. \begin{aligned} v^{-1} &= 4\pi \int_0^\infty dp p^2 z [\exp(\hbar^2 p^2 / 2mkT) - z]^{-1} \\ \beta P &= 4\pi \int_0^\infty dp p^2 \ln[1 - z \exp(-\hbar^2 p^2 / 2mkT)] \end{aligned} \right\}$$

which are known in classical analysis as Appell integrals, namely

$$(14) \quad \left. \begin{aligned} v^{-1} &= \lambda^{-3} g\left(\frac{3}{2}, z\right) \\ \beta P &= \lambda^{-3} g\left(\frac{5}{2}, z\right) \end{aligned} \right\} \quad \text{with} \quad \left\{ \begin{aligned} \lambda^2 &= 2\pi \hbar^2 / mkT \\ g(s, z) &= \frac{z}{\Gamma(s)} \int_0^\infty dt \frac{t^{s-1}}{\exp t - z} \end{aligned} \right. .$$

For every s with $Re(s) > 0$, g defines a function of z which is analytic in the cut complex plane $C \setminus [1, \infty)$. For $|z| < 1$ and $Re(s) > 0$, one receives the well-studied Lerch zeta functions which can be expanded in power series

$$(15) \quad g(s, z) = z \zeta(s, z) \quad \text{with} \quad \zeta(s, z) = \sum_{n=0}^{\infty} z^n (n+1)^{-s} \quad .$$

For $z = 1$ and $Re(s) > 1$ the above series converges to the Riemann zeta function $\zeta(s)$. Note that the values $s = \frac{3}{2}$ and $s = \frac{5}{2}$ — which are needed in (14) — fall within this range. Moreover $g(\frac{3}{2}, \cdot) : z \in (0, 1) \mapsto R^+$ is smooth, strictly increasing, with $\lim_{z \rightarrow 1} g(\frac{3}{2}, z) = \zeta(\frac{3}{2}) = 2.612 \dots$. The problem is thus mathematically under complete control.

Now to the physics. This divides into two steps.

The *first* step is easy: it considers the high temperature and low density regime, where $\lambda^3 v^{-1} < g(\frac{3}{2}, 1) = \zeta(\frac{3}{2})$. In particular, by straightforward 1st-order power expansion:

$$(16) \quad \text{for } \lambda^3 v^{-1} \ll 1 : P v = kT [1 - 2^{-5/2}(\lambda^3 v^{-1}) + \dots] \quad .$$

Hence, in this high temperature and low density regime, the quantum gas behaves asymptotically like the classical ideal gas of Boyle/Mariotte/Gay-Lussac. This is yet another confirmation of the Ehrenfests' remark according to which the classical limit obtains in QSP as a high temperature limit; note indeed that the so-called *thermal wavelength* λ that appears in (14) satisfies $\lambda \sim \hbar \beta^{1/2}$, i.e. in this problem again, the limits $T \rightarrow \infty (\Leftrightarrow \beta \rightarrow 0)$ and $\hbar \rightarrow 0$ have formally the same effect.

The *second* step in the treatment of the problem is where the bounty is to be found. The question is how to go beyond the above regime, i.e. beyond the unnatural limit

$$(17) \quad \lambda^3 v^{-1} = \zeta(\frac{3}{2}) \quad ,$$

a restriction no actual gas should be expected to respect. Mathematically, this limiting condition seems to appear as the consequence of the breakdown of analyticity in (14) that begins at $z = 1$. Physically, the problem appears because the limit $|\Lambda| \rightarrow \infty$ has been taken too carelessly.

Let us therefore return to the expression of v^{-1} when $|\Lambda| < \infty$. We have then, with $\langle n_k \rangle$ denoting the average number of particles in mode k :

$$(18) \quad \frac{1}{|\Lambda|} \sum_{k \in Z^3} \langle n_k \rangle = \frac{1}{|\Lambda|} \sum_{k \in Z^3, k \neq 0} \langle n_k \rangle + \frac{1}{|\Lambda|} \frac{z}{1-z} \quad .$$

As the $\langle n_k \rangle$ with $k \neq 0$ are well-behaved as $z \rightarrow 1$, the separation of (18) into two terms suggests that we take simultaneously the limits $|\Lambda| \rightarrow \infty$ and $z \rightarrow 1$ in such a manner that the second term in (18) approaches a finite limit, say v_o^{-1} , resulting in the replacement of (14) by:

$$(19) \quad \left. \begin{aligned} v^{-1} &= \lambda^{-3} \zeta(\frac{3}{2}) + v_o^{-1} \\ \beta P &= \lambda^{-3} \zeta(\frac{5}{2}) \end{aligned} \right\} \quad .$$

The above limiting procedure, interpreted as

$$(20) \quad v_o^{-1} = \lim_{|\Lambda|} \frac{1}{|\Lambda|} \langle n_o \rangle \quad ,$$

leads to a *macroscopic occupation of the ground state* $k = 0$; the theory does not predict the value of v_o : it may depend on the temperature. Note that the pressure P in (19) depends on temperature only (namely through λ). The state of the system described by (19) is called its *condensed phase*; the transition to this phase from the normal phase $\lambda^3 v^{-1} < \zeta(\frac{3}{2})$ is referred to as the *Bose–Einstein condensation*, (or *BEC*) and its appearance at low temperature is a prediction of purely quantum origin, one that has no equivalent in the classical world.

This begs for an instantiation in the world of the laboratory. At low temperature a superfluid phase appears in ${}^4\text{He}$. The density at the onset of this phenomenon is about $\rho \simeq .178 \text{ g/cm}^3$. Upon taking into account the value of the Avogadro number, one receives $v^{-1} \simeq 2.7 \cdot 10^{23} \text{ cm}^{-3}$, from which (17) gives a thermal wavelength $\lambda \simeq 4.6 \cdot 10^{-8} \text{ cm}$ which is not unreasonable for a quantity that is to be interpreted as a measure of the interparticle distance. To this value corresponds, via the definition of λ in (14), a temperature $T \simeq 3.2 \text{ K}$. The experimental value of the temperature at the onset of the superfluid phase in ${}^4\text{He}$ is $T \simeq 2.2 \text{ K}$, a rather remarkable fit, considering how crude the model is. Moreover, the thermodynamics of the model can be worked out — cf. e.g. [Huang, 1965] — and shows that the specific heat $C_v(T)$ at first increases monotonically from $C_v(0) = 0$ to exceed the classical value $3/2$ but then experiences a sharp peak — a discontinuity in the first derivative — from which it decreases monotonically to $\lim_{T \rightarrow \infty} C_v(T) = 3/2$. The specific heat of ${}^4\text{He}$ also exhibits such a singularity, albeit more pronounced: it is logarithmic; hence its name *λ – point*, as the graph of the specific heat as a function of temperature looks like the Greek lower case letter lambda.

All this represented a great success in the the mid-1920s. The next batch of problems appeared when the theory tried to account for the fact that ${}^4\text{He}$ is not a gas, but a liquid; for this, the *ideal* gas assumption of the model is quite unrealistic: a liquid is not made of non-interacting particles. Putting the interactions into the theory proved to be a formidable problem, long compounded by the experimental fact that ${}^4\text{He}$ was the only substance recognized to exhibit Bose–Einstein condensation: theoreticians had no variable parameter to guide and adjust their speculations. Following up on a proposal made in the late 1950s, the situation changed drastically during the 1980s and 1990s with the advent of micro-Kelvin technology which allowed BEC to be observed in atomic gases in harmonic traps; for two deep, but very different, reviews, cf. [Lieb, 2001] and [Pitaevskii and Stringari, 2003]; and for a brief overview [Emch and Liu, 2002, subsection 14.2.2].

The account in this subsection was limited mostly to the macroscopic, thermodynamical aspects of BEC in its infancy; in subsection 5.2 below, a C^* –algebraic treatment of the Bose–Einstein model is discussed in connection with the appearance in QSP of the modular structures to be associated to the equilibrium KMS condition.

2.5 Beyond the Bohr atom: the Thomas–Fermi model

The entry in quantum mechanics of Schrödinger wave-mechanics (1926) was marked by a resounding success: the physics community could recognize immediately the application of a then already well-established method to a new realm; the theoretical explanation of the energy spectrum of the hydrogen atom was reduced to solving an eigenvalue problem in a differential equation. Every entry-level text in quantum mechanics presents this derivation.

And yet, beyond the Bohr atom, the solution of the Schrödinger equation for an atom with even a few electrons turned to be an insurmountable task: the electrons are charged particles and while the interaction between a single electron and the nucleus had been rigorously accounted for in the hydrogen atom, one could not deal analytically with the mutual electromagnetic interactions between the electrons.

Very soon thereafter, Thomas [1927] and Fermi [1927] came up with a semi-classical model in which two ingredients enter. The first is the ground state electron density ρ which is assumed to be spherically symmetric and normalized by the condition

$$(21) \quad 4\pi \int_0^\infty dr r^2 \rho(r) = Z$$

where eZ is the charge of the nucleus. The second is the average electric potential $\Phi(r)$ in the atom. These two ingredients are assumed to satisfy the *classical* equation, the Poisson equation of electrostatics

$$(22) \quad \Delta\Phi \equiv \frac{1}{r} \frac{d^2}{dr^2}(r\Phi) = 4\pi e\rho \quad \text{with} \quad \lim_{r \rightarrow 0} \Phi(r) = eZ \quad .$$

And yet the model has a quantum aspect to account for the Pauli exclusion principle; this is the so-called Fermi–Dirac statistics that had been proposed just the previous year [Fermi, 1926]. Here, this shows up in:

$$(23) \quad n(r, p) = \begin{cases} 2h^{-3} & \text{if } \epsilon := \frac{1}{2m} - e\Phi < \epsilon_o \\ 0 & \text{if } \epsilon > \epsilon_o \end{cases} .$$

from which one gets, by integration over p (upon putting $\epsilon_o = 0$), that ρ satisfies

$$(24) \quad \rho(r) = \begin{cases} \frac{8\pi}{3h^3} (2me\Phi)^{3/2} & \text{if } \Phi > 0 \\ 0 & \text{if } \Phi < 0 \end{cases} .$$

Clearly, the model is conceptually inconsistent, with stakes in each of the classical and the quantum realms. Yet, in my student days this model was a routine staple of the quantum mechanics curriculum [Schiff, 1955; Landau and Lifshitz, 1958a; Messiah, 1960] as it can be solved without any further assumptions than those listed above; the solution is exact up to the fact that it requires a numerical computation well within the realm of a controllable approximation.

Upon using the numerical values of the Planck constant h , the charge e and the mass m of the electron, the model predicts that the radius of the atom, taken

to be the radius of the sphere that contains all the electrons but one increases monotonically from $2.2 \cdot 10^{-8}$ cm for $Z = 25$, to $2.8 \cdot 10^{-8}$ cm for $Z = 100$. The order of magnitude is correct. This can be counted therefore as an early success of quantum theory.

However, one should expect that such a crude model does not tell the whole story. Indeed: (1) the predicted increase stops at $Z = 55$ (corresponding to the cesium atom) after which the radius decreases, albeit slowly; (2) when looked at more closely, the model yields an electron density that has unreasonable properties both very close and very far from the nucleus. Besides, the model needs serious re-considerations to explain the existence of stable molecules or to accommodate a relativistic treatment. These problems never completely left the scene of theoretical physics, but remained somewhat in the background for about half-a-century, until rigorous analytic methods clarified the sense in which the model is asymptotically exact and may be used to study the stability of atoms, molecules and even stars; cf. [Lieb and Simon, 1977; Lieb, 1982a; Lieb, 1990]; see also [Catto *et al.*, 1998; Le Bris and Lions, 2005].

2.6 White dwarfs: the Chandrasekhar bound

Returning to the quantum ideal gas discussed in subsection 2.4, let us examine now the Fermi gas. Instead of (11), start with the partition function

$$(25) \quad Z(\Lambda, T, \mu) = \prod_{k \in Z^3} (1 + \exp^{-\beta(\epsilon_k - \mu)})$$

which now entails in the limit $\Lambda \uparrow Z^3$, instead of (13):

$$(26) \quad \left. \begin{aligned} v^{-1} &= 4\pi \int_0^\infty dp p^2 z [\exp(\hbar^2 p^2 / 2mkT) + z]^{-1} \\ \beta P &= 4\pi \int_0^\infty dp p^2 \ln[1 + z \exp(-\hbar^2 p^2 / 2mkT)] \end{aligned} \right\} .$$

In the high temperature and low density regime — $\lambda^3 v^{-1} \ll 1$ — one recovers again an asymptotic expansion, the leading term of which is the classical ideal gas:

$$(27) \quad \text{for } \lambda^3 v^{-1} \ll 1 : P v \simeq kT [1 + 2^{-5/2}(\lambda^3 v^{-1}) + \dots] .$$

Again, up to the sign of the correction, this is very similar to the Bose–Einstein result (16): it also coincides asymptotically with the classical ideal gas as T becomes large.

In the low temperature and high density regime — $\lambda^3 v^{-1} \gg 1$ — the situation differs drastically from what it was in subsection 2.4: whereas bosons tend to congregate, no two fermions are allowed in the same state on account of the Pauli exclusion principle. Recall that in chemistry, this is the principle that underpins a quantum explanation for the Mendeleev table of elements. In QSP the Pauli principle is visible through (26): in the ground state of the system, the fermions occupy the lowest possible energy states up to a finite energy, called the Fermi–energy

$$(28) \quad \epsilon_F = \frac{\hbar^2}{2m} [(3\pi^2)v^{-1}]^{\frac{2}{3}} .$$

For temperatures such that $kT \ll \epsilon_F$ the momentum distribution will be

$$(29) \quad \langle n_p \rangle = \begin{cases} 1 & \text{for } (|p|^2/2m) \lesssim \epsilon_F \\ 0 & \text{for } (|p|^2/2m) \gtrsim \epsilon_F \end{cases}$$

with a steep sigmoid of narrow breadth kT around ϵ_F . This regime is called the *degenerate Fermi gas*. To characterize this regime, rewrite $kT \ll \epsilon_F$, with ϵ_F as in (28), as:

$$(30) \quad \beta v^{-\frac{2}{3}} \gg \left[\frac{\hbar^2}{2m} (3\pi^2)^{\frac{2}{3}} \right]^{-1}$$

which gives a quantitative meaning to the expression *low temperature and high density* regime; for instance, this yields a useful first approximation for the gas of electrons in metals at usual temperatures. The condition $kT \ll \epsilon_F$ corresponds to $\lambda^3 v^{-1} \gg 1$ and in this regime (26) entails

$$(31) \quad Pv \simeq \frac{2}{5} \epsilon_F \left[1 + \frac{5\pi^2}{12} \left(\frac{kT}{\epsilon_F} \right)^2 + \dots \right] \quad \text{i.e.} \quad \lim_{\frac{kT}{\epsilon_F} \rightarrow 0} Pv^{\frac{5}{3}} = \frac{2}{5} (3\pi^2)^{\frac{2}{3}} \frac{\hbar^2}{2m} .$$

Hence, at fixed density, the pressure approaches a strictly positive constant as $T \rightarrow 0$, in marked contrast with the behaviour of the classical ideal gas (see (27)) where $T \rightarrow 0$ implies $P \rightarrow 0$.

Less mundane examples are provided by celestial objects, white dwarfs and neutron stars. With a temperature similar to that of the sun, i.e. $10^7 K$ to $10^8 K$ in the center, and a mass of the same order of magnitude as the sun, the white dwarfs have a very high density, about 10^6 to 10^7 times that of the sun. They are stars where all the hydrogen fuel has been burned, and thus they are constituted of completely ionized helium atoms. From these hypotheses on the composition and condition of a white dwarf, one computes the density of the electron gas, and then from (28) the corresponding Fermi energy ϵ_F which, when expressed in terms of $T_F = \epsilon_F/k$, turns out to give $T_F \simeq 10^{11} K$. Hence $T \ll T_F$ and it is consistent to assume — as R.H. Fowler did already in 1926 [Fowler, 1926] — that the electron assembly in the white dwarfs may be described as a degenerate Fermi gas, and that it is the enormous pressure in such a gas that prevents the star from gravitational collapse. It is however true that at such density and pressure, electrons must be treated relativistically, i.e. $\epsilon = \sqrt{(pc)^2 + (mc^2)^2}$ instead of $\epsilon = p^2/2m$. This brings about all sorts of analytic difficulties, among which is a change from $5/3$ towards $4/3$ in the power of v in (31). In the course of his computations of this effect, Chandrasekhar [1931a] remarked that since the gravitational pressure is governed by the mass of the star, the latter would collapse if the mass were to become too large; he actually evaluated this critical mass M_{max} to be

$$(32) \quad M_{max} \simeq (3\pi)^{\frac{1}{2}} \left(\frac{\hbar c}{G}\right)^{3/2} (\mu m_N)^{-2} \simeq 1.4 M_{\odot}$$

where (in cgs units) $\hbar = h/2\pi$ with $h \simeq 6.62 \times 10^{-27}$ ergs cm is the Planck constant, $c \simeq 3 \times 10^{10}$ cm/sec is the speed of light, $G \simeq 6.67 \times 10^{-8}$ dyn cm² g⁻² is Newton's gravitational constant, $m_N \simeq 1.66 \times 10^{-24}$ g, μ is the number of nucleons per electron; here $\mu = 2$ since the star is supposed to have used its hydrogen supply, and be made of ${}^4_2\text{He}$. Finally, to reduce the result in astronomic units, $M_{\odot} \simeq 1.99 \times 10^{33}$ g is the mass of the sun. Astronomers today refer to the maximum mass M_{max} as the *Chandrasekhar limit* [where mathematicians would speak of a "bound"].

Chandrasekhar's original derivation is mathematically correct, yet somewhat cumbersome. Already by the end of 1932, L.D. Landau [1932] presented a more elementary argument, *and* in addition, upon hearing of the discovery of the neutron, he applied the above formula to then putative neutron stars.

To have included these predictions here among the early "successes" of QSP may be justified only by hindsight. When they appeared in the early 1930s they and their consequences caused quite a wave, on the crest of which rode A.S. Eddington, an astronomer of commanding authority, who spoke of a *reductio ab absurdum* calling for the interposition of an as yet unknown fundamental theory: for him a massive star ($M > M_{max}$) collapsing to a black hole was heresy, and he was in a position not to mince his words about it. Eddington's fierce attack on a junior colleague did not cause Chandrasekhar to recant; unconvinced, Chandrasekhar nevertheless decided to turn to other astronomical problems until the late 1950s [Chandrasekhar, 1958] and early 1960s, when his speculations, and Landau's, found observational confirmations.

For the unfolding of the resolution of the Chandrasekhar–Eddington conflict, cf. e.g. [Shapiro and Teukolsky, 1983], the title of which already indicates the complete extent to which Chandrasekhar was ultimately vindicated. A pristine, yet non-technical presentation of the physics of the Chandrasekhar bound may be read in [Thorne, 1994, chapter 4]; the story of the neutron stars, albeit more involved is also told there in [Thorne, 1994, chapter 5]; for the fundamental technical support, cf. [Weinberg, 1972, chapter 11].

3 AXIOMATIC PRUNINGS

Usually, either one of two reasons prompts the process of axiomatization. The first is the search for the soul — some would say the skeleton — hidden inside the aleatory appearances of the body: a ritual of purification. The second is the need for fundamental changes when a theory faces increasingly insuperable limitations. Both of these reasons motivate the developments I retrace in the present section; as I write this I am reminded of the essential tensions described elsewhere in [Segal, 1990].

It is an interesting coincidence that the early 1930s saw the almost simultaneous — albeit independent — axiomatizations of two of the ingredients of QSP:

Quantum mechanics with von Neumann’s treatise [von Neumann, 1932c]; and *Statistics*, a.k.a. probability and stochastic processes, with Kolmogorov’s paper [Kolmogorov, 1933]. As both of these belong to other chapters of this Handbook, only a few words will suffice here.

3.1 Kolmogorov’s and von Neumann’s formalisms compared

In a nutshell, Kolmogorov’s syntax for probability starts with a seminal description of measure theory: a triple $\{\Omega, \mathcal{E}, \mu\}$ is given where \mathcal{E} is a σ -algebra of measurable subsets of a set Ω , and μ is a countably additive function

$$(33) \quad \mu : E \in \mathcal{E} \mapsto \mu(E) \in \mathbb{R}^+ \quad \text{with} \quad \mu(\Omega) = 1$$

i.e. μ is a *probability measure* on $\{\Omega, \mathcal{E}\}$. μ naturally extends to a functional on the algebra $\mathcal{A} = \mathcal{L}^\infty(\Omega, \mathcal{E}, \mu)$ of all essentially bounded functions $A : \Omega \rightarrow \mathbb{C}$:

$$(34) \quad \mu : A \in \mathcal{A} \mapsto \mu(A) = \iint_{\Omega} d\mu(\omega) A(\omega) \in \mathbb{C} .$$

Hereafter, I will refer to this extension as a *classical state*.

Similarly, von Neumann’s syntax involves a triple: $\{\mathcal{H}, \mathcal{P}, \psi\}$ where \mathcal{P} is the orthomodular lattice of all closed subspaces of a Hilbert space \mathcal{H} , ψ is a countably additive positive function

$$(35) \quad \left. \begin{aligned} \psi : P \in \mathcal{P} \mapsto \psi(P) \in \mathbb{R}^+ & \quad \text{with} \quad \psi(I) = 1 \quad \text{and} \\ \psi(\sum_n P_n) = \sum_n \psi(P_n) & \quad \forall \{P_n\} \subset \mathcal{P} \text{ such that } n \neq m \models P_n \perp P_m \end{aligned} \right\} .$$

I shall refer to any such function ψ as a quantum state. Gleason’s theorem asserts in particular — see below for a complete statement — that for every quantum state ψ there exists a density operator, i.e. a positive operator ρ of unit trace, such that ψ extends to the W^* -algebra $\mathcal{B} = \mathcal{B}(\mathcal{H})$ of all bounded linear operators from \mathcal{H} into itself:

$$(36) \quad \psi : B \in \mathcal{B} \mapsto \psi(B) = \text{Tr} \rho B \in \mathbb{R} .$$

When working within the von Neumann formalism, I will identify any closed subspace $P \subseteq \mathcal{H}$ and the projector $P \in \mathcal{B}(\mathcal{H})$ on this subspace; I will indifferently refer to ψ or to ρ as a *state* on \mathcal{B} ; and I will refer to the restriction of ψ to \mathcal{P} as a quantum measure. I will also follow the physicist’s custom of referring to ρ as a *density matrix*, thus ignoring the mathematician’s distinction between an operator and its expression in a specified (orthonormal) basis.

The mathematical similarities and differences between the classical and quantum realms are emphasized by the Koopman formalism of classical mechanics; cf. e.g. [Emch and Liu, 2002, pp. 255, 267]. This formalism — actually a precursor of the GNS construction — associates to $\{\Omega, \mathcal{E}, \mu\}$ the Hilbert space $\mathcal{H} = \mathcal{L}^2(\Omega, \mathcal{E}, \mu)$ of all functions $\Psi : \omega \in \Omega \rightarrow \Psi(\omega) \in \mathbb{C}$ that are square-integrable with respect to μ . Every element $A \in \mathcal{A} = \mathcal{L}^\infty(\Omega, \mathcal{E}, \mu)$ is then viewed as an element of $\mathcal{B} = \mathcal{B}(\mathcal{H})$,

namely under the identification of the function $A : \omega \in \Omega \mapsto A(\omega) \in \mathbb{C}$ with the multiplication operator $A : \Psi \in \mathcal{H} \mapsto A\Psi \in \mathcal{H}$ where $(A\Psi)(\omega) = A(\omega)\Psi(\omega)$. Under this identification \mathcal{A} becomes a maximal abelian W^* -subalgebra of \mathcal{B} ; while the center of \mathcal{B} , namely $\{C \in \mathcal{B} \mid \forall B \in \mathcal{B} : [B, C] = 0\}$ is trivial, i.e. consists of the multiples of the identity operator. Note further that every element $B \in \mathcal{B}(\mathcal{H})$ can be viewed as a continuous linear functional on the Banach space $\mathcal{T}(\mathcal{H})$ of all trace-class operators, spanned by the countably additive states; namely $B : T \in \mathcal{T}(\mathcal{H}) \mapsto \text{Tr} TB \in \mathbb{C}$; conversely every norm-continuous linear functional on $\mathcal{B}(\mathcal{H})$ obtains in this manner; i.e. $\mathcal{B}(\mathcal{H})$ is the Banach space *dual* of $\mathcal{T}(\mathcal{H})$; equivalently, $\mathcal{T}(\mathcal{H})$ is the *predual* of $\mathcal{B}(\mathcal{H})$. Similarly, the predual of $\mathcal{L}^\infty(\Omega, \mathcal{E}, \mu)$ is the Banach space of $\mathcal{L}^1(\Omega, \mathcal{E}, \mu)$, spanned by the probability distributions which are absolutely continuous with respect to μ .

The interpretation of a quantum state ψ in terms of classical probabilities obtains upon reading (35) separately for each family $\{P_n\}$ of mutually compatible quantum events. The bijective equivalence between the objects described by (35) and (36) is the pragmatic content of Gleason's theorem; cf. e.g. [Emch and Liu, 2002, p. 225]: every quantum state can be uniquely written in the form (36), and every density operator ρ defines through (36) a function ψ satisfying (35), i.e. a quantum state ψ . For the semantic, i.e. the empirical (frequentist vs. subjective) interpretations of states, first in classical probability theories, and then in quantum theories, cf. e.g. [Jaynes, 1967; Emch and Liu, 2002; Emch, 2005]; in particular, see [Uffink, 2006] for the evolution in CSP of the primacy of one over the other of these interpretations of probabilities.

Again in a nutshell, I believe that it serves my purpose well, in most of this essay, to espouse the 'subjective' rather than their 'frequentist' interpretation, namely *to view the state of a physical system* — be it classical or quantum, macroscopic or microscopic — *as a faithful summary of the knowledge one has of the process by which this system has been prepared*. In particular, this semantic view of the quantum state shall translate well from the case of systems with *finitely many degrees of freedom* considered in von Neumann's quantum mechanics, to the systems *with infinitely many degrees of freedom* to be considered in QSP; see subsections 3.4 to 6.3. In particular, while von Neumann's beams or 'ensembles', of independent, identically prepared systems — [von Neumann, 1932c, note 156] — are adequate to describe scattering experiments or the atomic spectroscopy of his time, the view of quantum states that I choose to adopt here accomodates better the description of single macroscopic systems — such as a cup of coffee or a measuring apparatus.

3.2 QSP in von Neumann's formalism

The centerpiece of equilibrium QSP in von Neumann's formalism is the following result [von Neumann, 1932c]:

THEOREM 3. *Let \mathcal{H} be a Hilbert space, H be a self-adjoint operator acting in \mathcal{H} and such that for all $\beta > 0$: the partition function $Z := \text{Tr} \exp(-\beta H)$ be finite. And, with $k > 0$ fixed, let for any state ρ on $\mathcal{B}(\mathcal{H})$*

$$(37) \quad S[\rho] = -k \operatorname{Tr} \rho \log \rho .$$

As H has discrete spectrum and is bounded below, let ϵ_o be its smallest eigenvalue; and let s denote either the largest eigenvalue of H if H is bounded above, or ∞ if it is not. Then, for any given $\epsilon_o < E < s$, the maximum of $S[\rho]$, subject to the constraint $\operatorname{Tr} \rho H = E$ is reached on the state

$$(38) \quad \rho = Z^{-1} e^{-\beta H} \quad \text{with} \quad Z = \operatorname{Tr} e^{-\beta H}$$

where the value of β is determined by the value E of the constraint.

The first part of the proof consists in showing that the maximum occurs on the class of states of the form $\rho = \sum_n \lambda_n P_n$ where $\sum_n \epsilon_n P_n$ is the spectral resolution of H . After this, the result follows from the classical argument using Lagrange multipliers with respect to the collection of variables $\Lambda = \{\lambda_n\} \subset \mathbb{R}^+$, namely from determining the maximum of the function $S[\Lambda] = -k \sum_n \lambda_n \log \lambda_n$ subject to the simultaneous constraints $\sum_n \lambda_n \epsilon_n = E$ and $\sum_n \lambda_n = 1$.

Note that this variational principle could have been rephrased as defining the state ρ in (38) as the state that minimizes — now under the single constraint $\operatorname{Tr} \rho = 1$, i.e. $\sum_n \lambda_n = 1$ — the Helmholtz free-energy defined as $F := E - TS$ with E and S as in the theorem, and $\beta = kT$ where k is known as the Boltzmann constant (see below).

Note also that, in either of these two forms, this variational principle has its root in the classical statistical physics (CSP) of Boltzmann and Gibbs; cf. [Uffink, 2006]. Conceptually, and very much as in CSP, the von Neumann QSP result involves a consensus on two questions. The first question is to justify the interpretation of S as an entropy. There are two ways to do this.

- (i) Firstly, as in CSP, one may identify S with the equilibrium entropy of macroscopic thermal physics upon computing S for well-controlled model(s), such as the ideal gas and finding — for in each of the specific cases considered — that the value of S_{max} obtained through the above theorem coincides with the value of the thermodynamical entropy. It is only at that stage that k may be identified with the universal Boltzmann constant $k \simeq 1.3810^{-23}$ J/deg; note the units, namely [energy]/[temperature], as is proper for the thermal entropy where T is the integrating factor that allows one to pass from the “heating” differential η to the exact differential $dS = \eta/T$. As fine as that may be for equilibrium CSP/QSP, this identification leaves open the interpretation of S as entropy in non-equilibrium situations.
- (ii) The second route to an interpretation of S is to show that $I(\rho) = -S(\rho)$ is a measure of the information content of the state ρ , namely to find empirically meaningful conditions that express the intuitive concept of “information content” and to show that — up to a multiplicative constant — there exists exactly one S that satisfies these conditions. The argument offered by Khinchin [1957] for classical probability distributions involves — *inter alia* — the axiom of consistency under refinements. This argument was transposed to the quantum case by Thirring [1983b] to give:

THEOREM 4. $S[\rho] = -k \operatorname{Tr} \rho \log \rho$ is the only functional satisfying:

1. $S[\rho]$ is continuous in ρ , in the sense that it is a continuous function of the eigenvalues of ρ .
2. For every finite probability distribution $P = \{p_n \mid n = 1, 2, \dots, N\}$ and every finite collection of states $\{\rho_n \mid n = 1, 2, \dots, N\}$ on a finite collection of Hilbert spaces $\{\mathcal{H}_n \mid n = 1, 2, \dots, N\}$, let ρ be the state defined on $\mathcal{H} = \bigoplus_{n=1}^N \mathcal{H}_n$ by $\rho = \bigoplus_{n=1}^N p_n \rho_n$. One has then: $S[\rho] = S[P] + \sum_{n=1}^N p_n S[\rho_n]$ where $S[P]$ is the value of the Khinchin functional for the probability distribution P .
3. $S \left[\begin{pmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{pmatrix} \right] = k \log 2$.

The first of the three conditions of theorem 4 is clear: an arbitrarily small change in the state should result in an arbitrarily small change in the information it conveys. The second condition expresses refinement under a particular class of partitionings; while the third is only a normalization. Just as in CSP, the *quantum information content* ($-S$), uniquely specified by these conditions, is formally used to define the *quantum entropy* S .

The second question concerning the conceptual relevance of theorem 3 is to justify the very use of a *variational principle*; compare with [Uffink, 1995]. In my view, for both the classical and the quantum cases, this comes most naturally when one opts for the subjective interpretation of states rather than the frequentist interpretation. Indeed, if one wants the state to account for the knowledge one has of the system, it seems consistent to select for ρ the state that assumes no more information than that expressed explicitly by the constraint.

When the operator H in theorem 3 is taken to represent the energy of the system, the state (38) is called — by analogy to the Gibbs canonical equilibrium state of CSP — the quantum canonical equilibrium state for the natural temperature $\beta = 1/kT$. Note in particular that, in the Schr odinger picture, the evolution generated by H , namely:

$$(39) \quad \forall t \in \mathbb{R} : \quad \rho(t) = U(t)\rho U(-t) \quad \text{with} \quad U(t) = \exp^{-i\frac{1}{\hbar}Ht}$$

leaves the canonical equilibrium state invariant, as is to be expected when one wishes to identify the energy-operator with the Hamiltonian of the system.

At first sight, the von Neumann formalism affords a good start for the development of a quantum ergodic theory. To keep things as simple as possible, consider the Hilbert space $\mathcal{L} = \{X \in \mathcal{B}(\mathcal{H}) \mid \operatorname{Tr} X^*X < \infty\}$ equipped with the scalar product $(X, Y) = \operatorname{Tr} X^*Y$. This space is known to mathematicians as the space of Hilbert–Schmidt operators acting on \mathcal{H} . In particular, every density matrix is an element of \mathcal{L} ; and thus this space is also known to physicists as the Liouville space of the quantum system described on \mathcal{H} . The advantage of restricting attention to this space is that (39) extends to a unitary action on \mathcal{L} :

$$(40) \quad V : (t, X) \in \mathbb{R} \times \mathcal{L} \mapsto V(t)[X] = U(t)XU(-t) \in \mathcal{L} \quad .$$

In the same way as the self-adjoint generator H of the continuous unitary group $\{U(t)|t \in \mathbb{R}\}$ is called the Hamiltonian of the quantum system considered, the self-adjoint generator L of the continuous unitary group $\{V(t)|t \in \mathbb{R}\}$ is called the *Liouillian* of this system. One has then

THEOREM 5. *Let $H \in \mathcal{B}$ have purely discrete spectrum, i.e. H can be written in the form $H = \sum_n \epsilon_n P_n$ where the P_n are mutually orthogonal projectors adding to I . Then the following limit exists*

$$(41) \quad E_{erg}[X] = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt V(t)[X] \quad \text{with } X \in \mathcal{L} \quad ;$$

and

$$E_{erg}[X] = \sum_n P_n X P_n \quad \text{and} \quad \forall t \in \mathbb{R} : V(t)[E_{erg}[X]] = E_{erg}[X].$$

In particular, the ergodic average $E_{erg}[\rho]$ of a density matrix ρ exists, is again a density matrix, and is time-invariant.

It is tempting to try and consider theorem 5 as proper quantum version of the classical ergodic theorems of Birkhoff [1931] or von Neumann [1932a]. Indeed, the *conclusions* of these classical theorems and of theorem 5 are similar when one reads them in terms of (countably additive) ‘states’ respectively defined as:

- $A \in \mathcal{L}^\infty(\Omega, \mu) \mapsto \int_\Omega d\mu f A \in \mathbb{C}$ where $f \in \mathcal{L}^1(\Omega, \mu)$, f positive with f normalized by $\int_\Omega d\mu f = 1$ (for the classical case);
- $A \in \mathcal{B}(\mathcal{H}) \mapsto \text{Tr } \rho A \in \mathbb{C}$ where ρ is a density matrix, i.e. a positive trace-class operator with ρ normalized by $\text{Tr } \rho = 1$ (for the quantum case);

and similarly for their respective time-averages.

Note that while the classical theorems are usually followed by a corollary involving the (quasi-)ergodic hypothesis and some discussion of the relevance of the results for the foundations of CSP — for a critical presentation see e.g. [Uffink, 2006, section 6.1] — I do not intend to try and follow suit here, in view of theorems 7 and 8 below which, for the purposes of QSP, cast a shadow on the adequacy of the *assumptions* theorem 5 makes on the Hamiltonian H . For a quantum ergodic theorem better adapted to the needs of QSP, see theorem 25 below.

Nevertheless, two related interesting comments may be made about theorem 5.

- (i) If, in this theorem, H is non-degenerate, i.e. if $\forall n : \dim P_n = 1$, then $E_{erg}[\rho]$ coincides with

$$(42) \quad Q_o[\rho] = \sum_n \text{Tr}(\rho P_n) P_n = \sum_n (\rho \Psi_n, \Psi_n) P_n$$

where $P_n \Psi_n = \Psi_n$ with $(\Psi_n, \Psi_m) = \delta_{mn}$, and where $Q_o[\rho]$ is thus the density matrix resulting from the von Neumann quantum measuring process [von

Neumann, 1932c, p. 351]; see also subsection 6.3 below. In particular, if ρ is a pure state, i.e. is a projector P_Ψ on some vector $\Psi = \sum_n c_n \Psi_n$, then $Q_o[P_\Psi] = \sum_n |c_n|^2 P_n$ has lost all the information encoded in the relative phases of the coefficients c_n .

- (ii) In [von Neumann, 1932c, pp. 380 ff] von Neumann shows that the entropy S of a state does not decrease — and in the generic case does increase — as the result of a measurement, whereas it is constant under the unitary evolution (40). He thus sees in

$$(43) \quad S[Q_o[\rho]] \geq S[\rho]$$

a confirmation that quantum measurements are generically *irreversible* processes. Similarly then, the information encoded in a (non-degenerate) density matrix ρ may only decrease as a result of taking its time-average, a reasonable feature indeed.

Yet, while theorem 5 could have been regarded as the germ of a quantum ergodic theory, the occurrence of monotonic irreversibility in QSP is significantly more elusive, as the next subsection demonstrates.

3.3 Some reasons to go beyond von Neumann’s formalism

Some of the problems non-equilibrium QSP has to face are illustrated in a simple spin-lattice model that was originally suggested to me by an actual experiment, the so-called nuclear free-induction relaxation; cf. [Emch and Liu, 2002, section 15.3].

The system consists of a linear chain of N interacting spins $\{\sigma_k = (\sigma_k^x, \sigma_k^y, \sigma_k^z) \mid k = 1, \dots, N\}$ with N even (and large, in a sense to be specified later on), and let

$$(44) \quad \sigma_k^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_k^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_k^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

be the Pauli matrices acting on $\mathcal{H}_k \simeq \mathbb{C}^2$. The Hilbert space of the system is then $\mathcal{H} = \otimes_k \mathcal{H}_k \simeq \mathbb{C}^{2^N}$. In this chain, two spins on sites k and $k + n$ interact with an energy $-J_n \sigma_k^z \sigma_{k+n}^z$, with $J_n > 0$ so that a lower energy is ascribed to configurations in which the z -components of these spins are parallel rather than anti-parallel. The whole system is plunged in a homogeneous magnetic field B in the direction z . The total Hamiltonian is chosen to be

$$(45) \quad H_N = -B \sum_{k=1}^N \sigma_k^z - \sum_{k=1}^N \sum_{n=1}^{N/2} J_n \sigma_k^z \sigma_{k+n}^z \quad \text{with} \quad J_n = 2^{-n} J_o > 0.$$

The system is initially prepared in the state

$$(46) \quad \rho_N = Z_N^{-1} \exp^{-\beta B \sum_{k=1}^N \sigma_k^z} \quad \text{with} \quad Z_N = \text{Tr} \exp^{-\beta B \sum_{k=1}^N \sigma_k^z}.$$

For the three “macroscopic” observables

$$(47) \quad S_N^\alpha = \frac{1}{N} \sum_{k=1}^N \sigma_k^\alpha \quad \text{with } \alpha \text{ standing for } x, y, z$$

one computes easily from (39–40) with $H = H_N$ given by (45):

$$(48) \quad \left. \begin{aligned} \text{Tr}(V_N(t)[\rho_N] S_N^x) &= \text{Tr}(\rho_N S_N^x) \cos(2Bt) f_N(t) \\ \text{Tr}(V_N(t)[\rho_N] S_N^y) &= \text{Tr}(\rho_N S_N^y) \sin(2Bt) f_N(t) \\ \text{Tr}(V_N(t)[\rho_N] S_N^z) &= \text{Tr}(\rho_N S_N^z) \end{aligned} \right\}$$

where

$$(49) \quad \left. \begin{aligned} f_N(t) &= f(t)/W_N(t) \quad \text{with} \\ f(t) &= \left[\frac{\sin(J_o t)}{J_o t} \right]^2 \quad \text{and} \quad W_N(t) = \left[\frac{\sin(2^{-N/2} J_o t)}{2^{-N/2} J_o t} \right]^2 \end{aligned} \right\} .$$

REMARKS 6.

1. For the purpose of discussing the putative irreversibility of the model, the (conservative) Larmor precession $\{\cos(Bt), \sin(Bt)\}$ of the magnetization around the direction z of the magnetic field B is of little or no interest.
2. In favour of the “irreversibility” of the model, one first notes that

$$(50) \quad \forall t \text{ with } |t| \ll T_N = 2^{N/2} \pi J_o^{-1} : \quad f_N(t) \simeq f(t)$$

and then the decay of $|\text{Tr} V_N(t)[\rho] S_N^\alpha|$ is governed by t^{-2} . Therefore, *in this time frame*, the magnetization (48) exhibits an apparent approach to equilibrium.

3. However, against the statement that the model would show an irreversible approach to equilibrium, one observes that

$$(51) \quad \lim_{t \rightarrow T_N} f_N(t) = 1 = f_N(0)$$

and thus, over the long run the system is periodic in time. This quantum model therefore would seem to confirm the classical Zermelo recurrence objection, or *Wiederkehrerwand*; for the latter, see [Uffink, 2006, section 4.5].

4. The saving grace, nevertheless, is that the period T_N increases exponentially with the size N of the system; see (50). This exponential behaviour is already encountered in CSP, as demonstrated by the Ehrenfest dog-flea model briefly mentioned in subsection 6.1 below. Thus, a modern Galileo would have his *Simplicio* argue that for macroscopically large systems, unaccountable perturbations would set in before T_N is approached, thus irremediably masking this periodicity; compare this to Boltzmann’s responses to the Zermelo objection; see again [Uffink, 2006, section 4.5].

5. Upon taking stock of this objection *Salviati* would invoke some modern version of the apocryphal commandment to the effect that “*Thou shalt not interchange limits*” since:

$$(52) \quad \lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} f_N(t) \text{ does not exist} \quad \text{but} \quad \lim_{t \rightarrow \infty} \lim_{N \rightarrow \infty} f_N(t) \text{ exists and is } 0 .$$

6. The present model presents also a quantum manifestation of the classical Loschmidt reversibility objection, or *Umkehrinwand* — cf. [Uffink, 2006, section 4.3] — as one has, here also:

$$(53) \quad f_N(-t) = f_N(t) \quad \text{and even} \quad f(-t) = f(t)$$

confirming the classical *Janus* dictum according to which the security of a *postdiction* is the same as that of a *prediction*. Hence, this model indicates that, if the *Umkehrinwand* were indeed a genuine objection to QSP — which I do not believe it is — the thermodynamical limit would not avoid it, whereas remarks (4) and (5) above show how it may respond to the *Wiederkehrinwand*.

7. Finally, one serious shortcoming of the present model — not as a model of the particular experiment considered above, but as a model for the approach to equilibrium in a QSP accounting for transport coefficients — is that, even in the limit of $N \rightarrow \infty$, the evolution is monitored by an inverse power law in time, rather than an exponential law, as would be required for the type of behaviour encountered in such macroscopic situations as described by Newton’s cooling law, Fourier’s heat equation, or more generally any macroscopic differential transport equation with linear coefficients.

The model illustrates explicitly some of the essential limitations of the von Neumann formalism for QSP, as manifested in the following two general results. The main assumption of theorems 3 and 5, namely that the Hamiltonian operator H has discrete spectrum, though seemingly innocuous — and actually necessary — when dealing with equilibrium QSP, has one potentially disastrous consequence when one attempts to extend the formalism to non-equilibrium situations: metastases of the classical objections spread into the quantum realm, as we shall now see.

The first result is a quantum version of the classical recurrence theorem of Zermelo. To be mathematically precise recall, in the words of Besicovitch’s standard text [Besicovitch, 1954], that a function $f : t \in \mathbb{R} \mapsto f(t) \in \mathbb{C}$ is said to be *almost periodic* in the sense of Harald Bohr, if $f(t + T)$ is approximately equal to $f(t)$ — with an arbitrary degree of accuracy — for infinitely many values of T , these values being spread over the whole real line, in such a way as not to leave empty intervals of arbitrarily great length.

THEOREM 7. *If the Hamiltonian $H = H^* \in \mathcal{B}(\mathcal{H})$ has purely discrete spectrum, i.e. if $H = \sum_n \epsilon_n P_n$; and if $\{V(t) \mid t \in \mathbb{R}\}$ is the unitary action in the Liouville space \mathcal{L} defined in (40), then*

$$(54) \quad \forall X, Y \in \mathcal{L} : f_{X,Y}(t) = \text{Tr}(V(t)[X]Y)$$

is an almost periodic function in t in the sense of H. Bohr.

Proof. $f(t)$ is a Fourier series $\sum_{n,m} a_{n,m} \exp^{-i\frac{1}{\hbar}(\epsilon_n - \epsilon_m)t}$ with $a_{n,m} = \text{Tr}(P_n X P_m y)$; by the Schwartz inequality in $\mathcal{L} : \sum_{n,m} |a_{n,m}|^2$ converges and thus — cf. [Besicovitch, 1954] — f is an almost periodic function of t in the sense of H. Bohr. ■

One might then attempt to get rid of recurrences by assuming — as is certainly allowed in the von Neumann formalism of quantum mechanics, provided $\dim \mathcal{H} = \infty$ — that the spectrum of the Hamiltonian is purely continuous. From the point of view of QSP, however, this cure would raise the following new difficulty, namely that ergodic states may not be countably additive, i.e. may not be representable by density matrices.

To describe this phenomenon, consider the Banach space $\mathcal{B} = \mathcal{B}(\mathcal{H})$ equipped with its usual operator norm; and denote by \mathcal{B}^* its dual, i.e. the Banach space of all continuous, linear functionals on \mathcal{B} . Then

$$(55) \quad \mathcal{B}^* = \mathcal{A}^* \oplus \mathcal{A}^\perp$$

where

- i. \mathcal{A} is the space of compact operators on \mathcal{H} , i.e. $\mathcal{A} = \{A \in \mathcal{B}(\mathcal{H}) \mid \Psi_n \rightharpoonup \Psi \Rightarrow A\Psi_n \rightarrow A\Psi\}$; here, \rightharpoonup and \rightarrow respectively denote weak- and strong-convergences in \mathcal{H} . When the $*$ -algebra \mathcal{A} is equipped with the operator norm it inherits from $\mathcal{B}(\mathcal{H})$, \mathcal{A} is closed in $\mathcal{B}(\mathcal{H})$ and thus is a Banach space on its own; in fact \mathcal{A} is the only non-trivial closed two-sided $*$ -ideal of \mathcal{B} .
- ii. For every $\varphi \in \mathcal{A}^*$, the dual of \mathcal{A} , there exists a unique trace-class operator $R \in \mathcal{T} = \{B \in \mathcal{B} \mid \text{Tr}(B^*B)^{\frac{1}{2}} < \infty\}$ such that $\forall A \in \mathcal{A} : \varphi(A) = \text{Tr}(RA)$. In particular, to every positive, continuous linear functional ψ on \mathcal{A} such that $\sup_{A \in \mathcal{A}, \|A\| \leq 1} \|\psi(A)\| = 1$ there corresponds a unique density matrix, and conversely.
- iii. $\mathcal{A}^\perp = \{\varphi \in \mathcal{B} \mid A \in \mathcal{A} \Rightarrow \varphi(A) = 0\}$.

Note that each of the inclusions $\mathcal{T} \subseteq \mathcal{L} \subseteq \mathcal{A} \subseteq \mathcal{B}$ is strict iff \mathcal{H} is infinite-dimensional, a condition that is required whenever one wants to avoid recurrences, since $\dim \mathcal{H} < \infty$ obviously entails that the spectrum of H is purely discrete, and then theorem 7 applies.

We can now make precise the above mentioned difficulty concerning the description of ergodic states within the context of countably additive states:

THEOREM 8. *Let $H \in \mathcal{B}(\mathcal{H})$ be the self adjoint generator of any strongly continuous unitary group $\{U(t) \mid t \in \mathbb{R}\}$ acting on \mathcal{H} ; and, with t running over \mathbb{R} , let $\rho \in \mathcal{T} \mapsto \rho(t) = U(t)\rho U(-t) \in \mathcal{T}$ describe the evolution of any density matrix ρ ; further, let ψ_t denote the corresponding (countably additive!) state $\psi_t : B \in \mathcal{B}(\mathcal{H}) \mapsto \psi_t(B) = \text{Tr}(\rho(t)B) \in \mathbb{C}$. Then, it follows that:*

a. For every compact observable $A \in \mathcal{A}$ the ergodic limit

$$(56) \quad \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T dt \psi_t(A)$$

exists and defines a positive linear functional $E_\infty[\psi]$ on \mathcal{A} .

b. If, moreover, the spectrum of H is purely continuous, then $E_\infty[\psi]$ cannot be extended to a countably additive state on $\mathcal{B}(\mathcal{H})$.

Proof. For the part (a), the economical strategy is to take advantage of two density theorems, namely: (i) when \mathcal{L} is equipped with its Hilbert-Schmidt norm, it contains \mathcal{T} as a dense subspace; and (ii) when \mathcal{A} is equipped with the operator norm it inherits from $\mathcal{B}(\mathcal{H})$, it contains \mathcal{L} as a dense subspace. Hence, one can uniquely lift the evolution from $\mathcal{T}(\mathcal{H})$ to a unique unitary evolution on the Hilbert space \mathcal{L} where one can use the classical ergodic theorem — [von Neumann, 1932a], or [Emch and Liu, 2002; Uffink, 2006] — to assert the existence of the ergodic limit for any pair $(X, Y) \in \mathcal{L} \times \mathcal{L}$, and in particular for any pair $(\rho, A) \in \mathcal{T} \times \mathcal{L}$. Upon recalling the duality $\mathcal{A} = \mathcal{T}^*$, the ergodic result is then extended by continuity from $\mathcal{T} \times \mathcal{L}$ to $(\rho, A) \in \mathcal{T} \times \mathcal{A}$.

To prove part (b), one notices that, on the one hand, this limit is given, for every $A \in \mathcal{A}$ by $E_o[\psi](A) = \text{Tr}(\sum_n (P_n \rho P_n A))$ where $\{P_n\}$ is the set of all the projectors corresponding to the discontinuous jumps in the spectral family of H . Hence, when H has continuous spectrum, this set is empty and thus $\forall A \in \mathcal{A} : E_o[\psi](A) = 0$. On the other hand, $E_o[\psi]$ certainly extends further than \mathcal{A} ; for instance the ergodic limit obviously exists for any $X \in \{H\}'$, i.e. for all bounded observables that are constants of the motion; in particular $E[\psi_o](I) = 1$. Hence, even if $E[\psi_o]$ could be extended to a state on \mathcal{B} , this state would belong to \mathcal{A}^\perp and thus would have no countably additive component in the direct sum decomposition (55). ■

The proof of the theorem shows that the same objection can be raised with any Hamiltonian the spectrum of which contains even only one interval of continuity. Taking Hamiltonians that are still self-adjoint, but not bounded above would only raise more technical problems without providing a solution to the basic limitation exposed in theorem 8.

Hence, von Neumann's formalism for QSP leads non-equilibrium QSP into the horns of a dilemma: either the evolution is almost periodic or the ergodic states are not countably additive. In particular, non-equilibrium states cannot approach asymptotic states that can be described by density matrices.

To make a bad situation even worse, Zeh discovered — admittedly, long after von Neumann's treatise had appeared and yet relevant to the thrust of this section — that there are serious empirical difficulties with the concept of an isolated quantum system [Zeh, 1970; Wigner, 1984]. Could certainty be fading out? [Prigogine, 1997]. Zeh's original observation has led to the development of the concept of

decoherence; cf. Landsman [Landsman, 2006, section 7.1]. I very briefly address this and some related issues in section 6 below.

Even in equilibrium QSP, the anchor provided by von Neumann was slipping: the formalism cannot account for the coexistence of thermodynamical phases; for a response to this objection, see subsection 5.7 below.

In counterpoint to these questions, one fundamental problem needs to be addressed: namely that the von Neumann formalism is not adequate to describe typical many-body systems where an infinite number of degrees of freedom are brought into the picture. The prescribed remedy is discussed in the next subsection.

3.4 Haag–Kastler’s axioms and Takeda’s inductive limits

This subsection outlines a formalism proposed to deal effectively with the non-relativistic many-body problems in QSP. This formalism was born out of the axiomatic responses prompted by the diagnosis of a mid-life crisis in relativistic Quantum Field Theory [QFT] [van Hove, 1952; Friedrichs, 1953; Wightman and Schweber, 1955]; the nail in the coffin was driven by Haag [1955]; cf. e.g. the famous Haag theorem and its embalming in [Barton, 1963, section 14], [Streater and Wightman, 1964, section IV.5], and/or [Emch, 1972a, section 3.d]. The algebraic axiomatization is presented here with sufficiently elementary details, yet with enough restraint to eschew the “imperialistic” label sometimes attached to it.

The main idea is to account for the local structure of infinitely extended systems. In their original proposal, Haag and Kastler [1964] mention several precedents in axiomatic QFT; among these [Haag, 1959a; Haag and Schroer, 1962]; see also [Haag, 1959b]. (I first heard of the algebraic approach in seminars in Geneva, where Araki presented some aspects of his Zurich lectures [Araki, 1961/2].) Segal’s early advocacy of an algebraic approach [Segal, 1947] ought also to be mentioned.

This subsection is divided in two complementary parts: the first part presents a description of the general structure; the second illustrates this structure with an example, the 1-dimensional quantum spin-lattice.

Part I. The general structure.

One begins by selecting an absorbing directed net \mathcal{F} of regions Λ of finite extension in space; usually, the space is the Minkowski space \mathbb{M}^{n+1} for relativistic QFT, the Euclidean space \mathbb{R}^n or a lattice \mathbb{Z}^n for non-relativistic QSP. The case of immediate physical interest is $n = 3$, but exploratory models are often constructed with $n = 1, 2$. Recall that a directed net is a partially ordered set — here the order relation is the usual set-theoretical inclusion — such that for every pair of elements Λ_1, Λ_2 in \mathcal{F} there is at least one element $\Lambda \in \mathcal{F}$ such that $\Lambda_1 \subseteq \Lambda$ and $\Lambda_2 \subseteq \Lambda$. To say that this net is absorbing is to say that for every point x in space there exists at least one element $\Lambda \in \mathcal{F}$ such that $x \in \Lambda$. The symbol $\Lambda_1 \bowtie \Lambda_2$ will be used to signify that two regions Λ_1 and Λ_2 are causally disjoint, i.e. in QFT, these

regions are spacelike to one another; and in non-relativistic QSP, they are disjoint in the set-theoretical sense, i.e. $\Lambda_1 \cap \Lambda_2 = \emptyset$. G denotes a group of rigid motions in the space, namely the inhomogeneous Lorentz group for \mathbb{M}^{n+1} ; the Euclidean group for \mathbb{R}^n ; or the group of lattice translations for \mathbb{Z}^n .

Secondly, to every $\Lambda \in \mathcal{F}$ one assigns a C^* -algebra \mathcal{A}_Λ ; without loss of generality, one may assume that \mathcal{A}_Λ has an identity I_Λ . This assignment is subject to the following three postulates.

POSTULATE 9 (Isotony). Whenever $\Lambda_1 \in \mathcal{F}$ and $\Lambda_2 \in \mathcal{F}$ satisfy $\Lambda_1 \subseteq \Lambda_2$, one is given an injective $*$ -homomorphism $i_{21} : \mathcal{A}_{\Lambda_1} \rightarrow \mathcal{A}_{\Lambda_2}$ such that

1. $i_{21}(I_{\Lambda_1}) = I_{\Lambda_2}$
2. $\Lambda_1 \subseteq \Lambda_2 \subseteq \Lambda_3 \Rightarrow i_{32} \circ i_{21} = i_{31}$.

The following result was proven by Takeda [1955].

THEOREM 10. Let \mathcal{F} be a directed net, and $\{\mathcal{A}_\Lambda \mid \Lambda \in \mathcal{F}\}$ satisfy the isotony postulate. Then there exist: a C^* -algebra \mathcal{A} with unit I ; and a family of injective $*$ -homomorphisms $\{i_\Lambda : \mathcal{A}_\Lambda \rightarrow \mathcal{A} \mid \Lambda \in \mathcal{F}\}$ such that

1. $\forall \Lambda \in \mathcal{F} : i_\Lambda(I_\Lambda) = I$;
2. $\Lambda_1 \subseteq \Lambda_2 \Rightarrow i_{\Lambda_1}(\mathcal{A}_{\Lambda_1}) \subseteq i_{\Lambda_2}(\mathcal{A}_{\Lambda_2})$;
3. $\bigcup_{\Lambda \in \mathcal{F}} i_\Lambda(\mathcal{A}_\Lambda)$ is a norm-dense sub- $*$ algebra of \mathcal{A} .

The C^* -algebra \mathcal{A} is called the C^* -inductive limit of the net $\{\mathcal{A}_\Lambda \mid \Lambda \in \mathcal{F}\}$. We will use hereafter the notations

$$\mathcal{A}_o := \bigcup_{\Lambda \in \mathcal{F}} i_\Lambda(\mathcal{A}_\Lambda) \quad \text{and} \quad \mathcal{A} = \overline{\mathcal{A}_o}.$$

POSTULATE 11 (Local commutativity). Whenever $\Lambda_1, \Lambda_2 \in \mathcal{F}$ satisfy $\Lambda_1 \bowtie \Lambda_2$, and $\Lambda_3 \in \mathcal{F}$ is such that both $\Lambda_1 \subseteq \Lambda_3$ and $\Lambda_2 \subseteq \Lambda_3$:

$$A_1 \in \mathcal{A}_{\Lambda_1} \text{ and } A_2 \in \mathcal{A}_{\Lambda_2} \Rightarrow i_{31}(A_1) i_{32}(A_2) = i_{32}(A_2) i_{31}(A_1) \quad .$$

The following result is then immediate.

COROLLARY 12. If $\Lambda_1 \bowtie \Lambda_2$, then

$$A_1 \in \mathcal{A}_{\Lambda_1} \text{ and } A_2 \in \mathcal{A}_{\Lambda_2} \Rightarrow i_{\Lambda_1}(A_1) i_{\Lambda_2}(A_2) = i_{\Lambda_2}(A_2) i_{\Lambda_1}(A_1) \quad .$$

For the aspects of QSP considered here, it will be an innocent abuse of language to refer to the above postulate as simply the *postulate of locality*.

POSTULATE 13 (Covariance). An action $\nu : (g, A) \in G \times \mathcal{A}_o \rightarrow \nu_g[A] \in \mathcal{A}_o$ is given so that for every region $\Lambda \in \mathcal{F}$, ν_g induces a $*$ -isomorphism between \mathcal{A}_Λ and $\mathcal{A}_{g[\Lambda]}$, where $g[\Lambda]$ denotes the image of the region Λ under the point transformation g .

Upon using theorem 10, this can be lifted to \mathcal{A} , namely $\nu_g[i_\Lambda(\mathcal{A}_\Lambda)] = i_{g[\Lambda]}(\mathcal{A}_{g[\Lambda]}) :$

COROLLARY 14. *The action of G extends by continuity to a norm-continuous group representation $\nu : g \in G \rightarrow \text{Aut}(\mathcal{A})$.*

DEFINITION 15. With the above notations, \mathcal{A}_o is called the algebra of local observables; and \mathcal{A} is called the algebra of quasi-local observables. Moreover, if φ is a state on \mathcal{A} such that $\forall g \in G : \varphi \circ \nu_g = \varphi$, let π_φ be the corresponding GNS representation. The von Neumann algebra $\mathcal{N}_\varphi = \pi_\varphi(\mathcal{A})''$ is called the algebra of global observables relative to the state φ .

Note that quasi-local observables involve norm limits; they are therefore general, algebraic objects that can be defined abstractly, i.e. without reference to any particular Hilbert space representation. In contrast, global observables that are not quasi-local involve weak-operator limits, and thus depend on the Hilbert space representation in which these limits are taken; for the purposes of QSP these observables depend, via the GNS construction, on the physical situation for which they are defined, i.e. on the state with respect to which they are considered. This aspect of the theory will be discussed in details in subsection 3.5 — see in particular the preliminaries to scholium 23, and remark 26(1) — and it will be essential for the treatment of phase transitions, inasmuch as these involve averages of observables, for instance the spontaneous magnetization in ferromagnets; see subsection 5.7 below.

Part II. A concrete example of a net of observable-algebras.

This example exhibits the construction of the algebra of observables for an infinite quantum spin-lattice system that obtains in the thermodynamical limit of finite systems such as the one treated in subsection 3.3 above. Consider indeed an infinite 1-dimensional lattice \mathbb{Z} with a quantum $\frac{1}{2}$ -spin sitting at each node (or “site”); hence a copy \mathcal{A}_k of the C^* -algebra $\mathcal{M}(2, \mathbb{C})$ of 2×2 matrices with complex entries is associated to each site $k \in \mathbb{Z}$; i.e. \mathcal{A}_k is generated by the Pauli matrices (44), i.e. by the three observables corresponding to the three components of a $\frac{1}{2}$ -spin sitting at site k .

Let now \mathcal{F} be the net of all finite subsets $\Lambda \subset \mathbb{Z}$. To each of these Λ is then associated the “local” C^* -algebra $\mathcal{A}_\Lambda = \otimes_{k \in \Lambda} \mathcal{A}_k$ which is thus a copy of $\mathcal{M}(2^{|\Lambda|}, \mathbb{C})$, where $|\Lambda|$ denotes the number of sites in Λ .

Let now Λ_1 and Λ_2 be two finite regions, with $\Lambda_1 \subseteq \Lambda_2$. An injective $*$ -homomorphism of \mathcal{A}_{Λ_1} into \mathcal{A}_{Λ_2} obtains by linearity from its restriction to monomials, namely

$$i_{21}(A_1 \otimes A_2 \otimes \dots \otimes A_{|\Lambda_1|}) = B_1 \otimes B_2 \otimes \dots \otimes B_{|\Lambda_2|}$$

with

$$\forall k \in \Lambda_2 : B_k = \begin{cases} A_k & \text{if } k \in \Lambda_1 \\ I_k & \text{if } k \notin \Lambda_1 \end{cases} .$$

These inclusions satisfy postulate 9.

Here, two finite regions are in the relation $\Lambda_1 \bowtie \Lambda_2$ whenever $\Lambda_1 \cap \Lambda_2 = \emptyset$. Since the commutators of observables attached to individual sites vanish whenever the two sites are different, any two observables attached to disjoint regions do commute. Formally, this is to say that postulate 11 is satisfied.

Finally, let $G := \mathbb{Z}$ denote the additive group of translations of the lattice. To define the action of G on the algebra of local observables it is sufficient to notice that for all $g \in G$ and all $\Lambda \in \mathcal{F} : |g[\Lambda]| = |\Lambda|$, so that both $\mathcal{A}_{g[\Lambda]}$ and \mathcal{A}_Λ are copies of the same matrix algebra, namely $\mathcal{M}(2^{|\Lambda|}, \mathbb{C})$: the images of a local observable and its translate are simply different copies of the same matrix; this indeed defines ν_g in such a way that postulate 13 is satisfied.

3.5 Quantum ergodic theory and macroscopic observables

While classical ergodic theory concerns itself with *measures* μ that are invariant under a group G and their *mixing* properties, quantum ergodic theory discusses the properties of G -invariant *states* and their *clustering* properties. Accordingly, in this subsection I will discuss also the roles of *space and/or time averages* in explaining at least part of the success of QSP; compare with [Uffink, 2006] and, in particular, with [Earman and R edei, 1996].

Therefore, one question to be addressed in this subsection must be whether and how ergodic theory may serve as a cornerstone to build up statistical mechanics. Traditionally, under the impetus of the emphasis the Ehrenfests' placed on Boltzmann's ergodic hypothesis (or rather its measure-theoretical version, the quasi-ergodic hypothesis), the group G is taken to be the group \mathbb{R} governing the *time evolution*. Nevertheless, partly in response to some swaying to and fro in Boltzmann's own writings, the jury is still out on the fundamental issue of this hypothesis' relevance for the foundations of CSP; see in particular [Uffink, 2006, section I.3, and subsections I.4.3 and I.6.1]. It is also remarkable that already Gibbs [Gibbs, 1902] chose to emphasize the role of mixing properties, i.e. properties which are stronger than metric transitivity and make more precise the presumption that the dynamics is 'erratic'; cf. e.g. [Uffink, 2006, sections I.4.1 and I.5]; other issues are touched upon in [Emch and Liu, 2002, pp. 317–330]; and, for a pertinent account that takes stock of the work done in the second half of the twentieth century, by the Lebowitz and Sinai schools, see [Szasz, 1996].

Consequently, I will concentrate here on two sub-questions: (i) the extent to which the mathematics of classical ergodic theory may be generalized to the formalism of quantum theory; and (ii) the extent to which such generalizations may help formulate better certain aspects of the foundations of QSP.

The answer to the first of these sub-questions is that much of the mathematics goes through, with some minor adjustments. The answers to the second is more

complex. On the one hand, as long as the focus remains on the time evolution the main issues persist, among which is the paucity of realistic models. On the other hand, when the group G deals with the geometry of the problem, quantum ergodic theory — especially and the roles of averages, and theorems linking extremal invariance and clustering properties — does help distinguish, within QSP, the quantum aspects of the microscopic description and the classical aspects of the macroscopic world. Thus, I divide the presentation in two parts, according to whether ergodicity is considered with respect to the time evolution or with respect to space symmetries.

A. Ergodicity with respect to time

Some insight is gained from a model first proposed in [Ford *et al.*, 1965], which comes in two versions, classical and quantum. The quantum version has been controlled mathematically in [Davies, 1972]. It is proven there that an infinite 1-dimensional chain of weakly coupled 1-dimensional quantum harmonic oscillators may serve as a thermal reservoir for a single 1-dimensional quantum oscillator in the chain and that a diffusion equation governs the evolution of the latter. This is accomplished by a rigorous treatment of the van Hove weak-coupling/long-time limit about which more will be said in subsection 6.1 below. Anticipating some mathematical definitions to be made precise later on — see paragraph 5.3.C — it is sufficient for the present discussion to register that, in this van Hove limit, a reduced evolution obtains which is a contractive semi-group of completely positive maps $\{\gamma_s \mid s \in \mathbb{R}^+\}$ acting on the von Neumann algebra $\mathcal{N}_o \simeq \mathcal{B}(\mathcal{L}^2(\mathbb{R}, \mathbb{C}))$ attached to the site of the single oscillator considered. Moreover this evolution, when observed from any one-dimensional subspace $\{x\mathbf{u} \mid x \in \mathbb{R}\}$ in the 2-dimensional phase space $\{\xi P + \eta Q \mid \zeta = (\xi, \eta) \in \mathbb{R}^2\}$ of the single oscillator, is described by a classical distribution $\mu(x, s)$ that satisfies for all $s \in \mathbb{R}^+$ the *diffusion equation*:

$$(57) \quad \partial_s \mu(x, s) = D \left[\partial_x^2 + \beta(V'(x) \partial_x + V''(x)) \right] \mu(x, s)$$

where $\beta = 1/kT$ is the natural temperature, $V = \frac{1}{2}\omega x^2$ is a harmonic potential, while the diffusion constant D and the frequency ω are numbers, the values of which depend only on the direction $\zeta/|\zeta| \in \mathbb{R}^2/S^1$. Note that the corresponding invariant measure is the canonical equilibrium, Gaussian measure $\mu(x) = Z^{-1} \exp(-\beta V(x))$ with $Z = \int_{\mathbb{R}} dx \mu(x)$, i.e. $Z^{-1} = \sqrt{2\pi\beta\omega}$.

The point of the model here is that the dissipative system described by the contractive *semi-group* $\{\gamma_s \mid s \in \mathbb{R}^+\}$ governing this Markovian diffusion process admits a canonical dilation to a conservative dynamical system. Indeed, there exists a *group* $\{\alpha_s \mid s \in \mathbb{R}\}$ of automorphisms of the von Neumann algebra $\mathcal{N} = \pi_\varphi(\mathcal{A})''$ describing the full chain of oscillators in the equilibrium state φ corresponding to the temperature β when the interactions are switched off. In conformity with subsection 3.4 the algebra of quasi-local observables \mathcal{A} is here the C^* -algebra $\otimes_{k \in \mathbb{Z}} \mathcal{N}_k$ where the \mathcal{N}_k are copies of \mathcal{N}_o . The free equilibrium state has the form $\varphi = \otimes \varphi_k$ where φ_k is the von Neumann canonical equilibrium state for the oscillator at the site k . Let now i be the injection of \mathcal{N}_o into \mathcal{N} and φ_o

denote the restriction of φ to \mathcal{N}_o , i.e. $\forall N_o \in \mathcal{N}_o : \varphi_o(N_o) = \varphi(i[N_o])$. Let further $E : \mathcal{N} \rightarrow \mathcal{N}_o$ be the canonical conditional expectation with respect to the state φ , satisfying $\varphi_o \circ E = \varphi$. The sense in which $\{\mathcal{N}, \alpha, E\}$ is a dilation of $\{\mathcal{N}_o, \gamma\}$ is that

$$(58) \quad \forall (s, N_o) \in \mathbb{R}^+ \times \mathcal{N}_o : \gamma_s[N_o] = E \circ \alpha_s \circ i[N_o] \quad .$$

For details, see [Emch, 1976] where, in particular, this result was noted to be very reminiscent of the classical flow of Brownian motion constructed by Hida [Hida, 1970] who also proved that this flow is a classical Kolmogorov flow, in the sense of the following definition.

DEFINITION 16. A classical dynamical system $\{\Omega, \mathcal{E}, \mu, \alpha^*\}$ consisting of a probability space $\{\Omega, \mathcal{E}\}$, a probability measure μ , and a group $\{\alpha^*_t \mid t \in \mathbb{R}\}$ of automorphisms of $\{\Omega, \mathcal{E}\}$ such that $\forall t \in \mathbb{R} : \mu \circ \alpha^*_t = \mu$, is said to be a classical Kolmogorov flow whenever there exists a σ -subring $\mathcal{A} \subset \mathcal{E}$ such that, with the notation $\mathcal{A}_t = \alpha^*_t[\mathcal{A}]$:

$$(1) \quad \forall t > 0 : \mathcal{A} \subset \mathcal{A}_t ; \quad (2) \quad \bigvee_{t \in \mathbb{R}} \mathcal{A}_t = \mathcal{E} ; \quad \text{and} \quad (3) \quad \bigwedge_{t \in \mathbb{R}} \mathcal{A}_t = \{\emptyset, \Omega\} .$$

Kolmogorov flows are characterized among classical dynamical systems by their having strictly positive dynamical entropy; thus they sit pretty high in the classical ergodic hierarchy, above the Lebesgue spectrum condition, and thus above the weaker conditions of mixing and ergodicity; for didactic accounts, cf. e.g. [Arnold and Avez, 1968; Cornfeld *et al.*, 1982].

The conservative quantum dynamical system described above as the canonical dilation of a contractive semigroup, does satisfy a quantum generalization of definition 16, namely:

DEFINITION 17. A quantum dynamical system $\{\mathcal{N}, \varphi, \alpha\}$ consisting of a von Neumann algebra \mathcal{N} , a faithful normal state φ on \mathcal{N} , and a group $\alpha = \{\alpha_t \mid t \in \mathbb{R}\}$ of automorphisms of \mathcal{N} , with $\forall t \in \mathbb{R} : \varphi \circ \alpha_t = \varphi$, is said to be a generalized Kolmogorov flow whenever there exists a von Neumann subalgebra $\mathcal{A} \subset \mathcal{N}$ such that, with the notation $\mathcal{A}_t = \alpha_t[\mathcal{A}]$:

$$(1) \quad \forall t > 0 : \mathcal{A} \subset \mathcal{A}_t ; \quad (2) \quad \bigvee_{t \in \mathbb{R}} \mathcal{A}_t = \mathcal{N} ; \quad (3) \quad \bigwedge_{t \in \mathbb{R}} \mathcal{A}_t = \mathbb{C}I ; \quad \text{and} \\ (4) \quad \forall t \in \mathbb{R} : \tau_t[\mathcal{A}] = \mathcal{A} ,$$

where $\{\tau_t \mid t \in \mathbb{R}\}$ is the modular group canonically associated to φ .

REMARKS 18.

1. The \bigvee in condition (2) involves a weak-operator closure, namely (2) means that \mathcal{N} is the smallest von Neumann algebra that contains all the \mathcal{A}_t ; the \bigwedge in condition (3) is simply the usual intersection; thus (3) signifies that no operator belongs to all \mathcal{A}_t unless it is a multiple of the identity.
2. The modular group τ will be introduced in section 4; let it suffice to say here that, if we were dealing with a finite system, τ would be the group of automorphisms of \mathcal{N} associated to the Hamiltonian corresponding to von Neumann's canonical equilibrium density matrix.

3. Definition 17 encompasses definition 16 when \mathcal{N} is taken to be the abelian von Neumann algebra $\mathcal{L}^\infty(\Omega, \mathcal{E})$ acting on the Hilbert space $\mathcal{H} = \mathcal{L}^2(\Omega, \mathcal{E}, \mu)$; in this case $\forall t \in \mathbb{R} : \tau_t = id$, and condition (4) is then trivially satisfied.
4. In the general case, condition (4) is necessary to ensure the existence of a conditional expectation $E : \mathcal{N} \rightarrow \mathcal{A}$.
5. Except for the positivity of the dynamical entropy — which depends on a consensus that is still pending about a physically meaningful definition of quantum dynamical entropy; see nevertheless [Narnhofer and Thirring, 1994b; Tuyls, 1998] and references therein — all the ergodic properties of classical Kolmogorov systems carry over straightforwardly from the classical to quantum realm [Emch, 1976]. In the model described above these properties are exhibited in the quantum triple $\{\mathcal{N}, \varphi, \alpha\}$.
6. Definition 17 was first proposed in [Emch, 1976]. Generalizations of this definition, involving the passage from W^* - to C^* -algebras, were then explored in [Narnhofer and Thirring, 1989].
7. The material of the present remark may be found in [Arnold and Avez, 1968] and is inserted here only as a preparation for the next remark. In classical ergodic theory the next rung up the ergodic ladder, just above Kolmogorov flows, is occupied by Anosov flows. These flows formalize an observation made in 1898 by Hadamard, namely that the geodesics on manifolds of negative curvature exhibit exponential sensitivity to initial conditions, in contrast with the usual linear sensitivity characteristic of free flows on flat manifolds. If the manifold is furthermore compact, one may intuitively expect that Hadamard's observation entails some kind of mixing behaviour. This is indeed the case: the first ever Hamiltonian flow shown to be ergodic — the geodesic flow on a compact surface of constant negative curvature — is already an Anosov flow. These flows exhibit exponentially contracting and expanding directions transversal to the direction of the flow, thus prefiguring a microscopic explanation for the empirically observed Lyapunov coefficients. The discrete-time archetype is the Arnold CAT map operating on the torus $T^2 := \mathbb{R}^2/\mathbb{Z}^2$. One ought to note that up to Kolmogorov flows, classical ergodic theory may be viewed as a chapter in probability theory; Anosov flows, in addition, involve an essential appeal to differential geometry, as was recognized only in the second half of the twentieth century through the work of the Russian school.
8. In order to explore possible quantum extensions of the concept of Anosov flow, a quantum analog of the latter has been devised by the present author in collaboration with Narnhofer, Sewell and Thirring [Emch *et al.*, 1994a]; for an antecedent, see [Benatti *et al.*, 1991a]; for a discussion of dynamical entropy in this context, see [Andries *et al.*, 1995]; for reviews and some general perspectives, see [Narnhofer, 2001; Narnhofer, 2005].

One essential feature of this extension is that now the phase space of this quantum CAT map is the noncommutative torus T_θ^2 , an ubiquitous staple of Connes' noncommutative geometry; cf. e.g. [Connes, 2000, section XIII] or [Garcia-Bondia *et al.*, 2003, chapter 12]; and for the place these tori occupy in the geometric quantization programme, cf. [Emch, 1998b]. As for quantum ergodic theory, it was noted already in [Emch *et al.*, 1994a] that the generators of the expanding and contracting horocycles form a basis in the 2-dimensional distinguished space of derivations that are not approximately inner — i.e. cannot be uniformly approximated by inner derivations [Garcia-Bondia *et al.*, 2003, section 12.3].

The presence of expanding and contracting directions in quantum as well as in classical Anosov flows offers a bridge from classical to quantum chaos. The problem of what is quantum chaos — or what it ought to be — has received attention from different perspectives; cf. e.g. [Gutzwiller, 1990]; for a philosophical perspective, cf. [Belot and Earman, 1997], and for a recent review, cf. in this volume [Landsman, 2006, section 5.6].

The investigations sketched in this remark, with applications to QSP in view, also have a mathematical parallel in QFT, cf. [Borchers, 1999; Wiesbrock, 1997]; see also subsection 5.5 below.

Summary and warning. It seems fair to infer that the *mathematical* generalization of classical ergodic kinematics to the quantum realm will carry through quite well. Nevertheless, the discussion of the underlying *physical* dynamics, when confronted with Hamiltonian mechanics, does not fare any more smoothly in the quantum case than it does in the classical case. Some of the conceptual problems may already be illustrated with the help of the model discussed at the beginning of this subsection. There, the dissipative dynamical system $\{\mathcal{N}_o, \gamma\}$ may be viewed as the reduced dynamics of two different conservative dynamical systems; both of these act on the same infinite assembly of harmonic oscillators. From the first system the reduced dynamics obtains only through the van Hove limit which compounds the very long-time effects — on a single subsystem — of a very weak coupling with, and within, the bath. But there is nothing in common between the time scale of the dynamics that governs the original conservative system and the time scale pertaining to the other conservative system, viz. the one obtained as the canonical dilation of the dissipative system. So there is little reason to believe that the ergodic behaviour of the latter reflects any global dynamical property of the former.

While this may be blamed on some naive modeling, it nevertheless emphasizes that the time scale of the conservative microscopic description and that of the emerging macroscopic description may differ significantly. In more sophisticated models, this will have to be taken into consideration and the complicated behaviour of the microscopic description may have to be washed away — one way or another — before a clean ergodic behaviour is manifested at the macroscopic level. It appears that van Hove's idea is a reasonable way to do this; see subsection 6.1

below.

Starting with their initial motivation in Boltzmann’s works, most presentations of classical ergodic theory focus on the properties of the time-evolution, in particular on the transitivity of measures and the *time*-averages of observables. Its generalization to the quantum realm invites the consideration of other aspects of classical ergodic theory, namely the *space* averages with respect to the actions of *other* groups beside those that govern the evolution. This will be done in the second part of this subsection.

B. Ergodicity with respect to space

As was already recognized by Haag [1959b] for QFT, the “other” group of most immediate relevance to QSP is the group of space translations, introduced as a part of the postulate of covariance in the Haag–Kastler axioms; cf. postulate 13 above. With $n = 1, 2, \dots$, let \mathbb{X}^n denote either the Euclidean space \mathbb{R}^n or the “cubic” lattice \mathbb{Z}^n ; and let $|x|$ denote the length of the vector $x \in \mathbb{X}^n$. Henceforth, we concentrate on the abelian group $G \simeq \mathbb{X}^n$ of all translations $x \in \mathbb{X}^n \mapsto x + a \in \mathbb{X}^n$ where $a \in \mathbb{X}^n$. Let further $\{\mathcal{A}_\Lambda \mid \Lambda \in \mathcal{F}\}$ be the corresponding Haag–Kastler net of local algebras, and \mathcal{A} be their C^* -inductive limit, with \mathcal{A} equipped with the group of automorphisms $\{\nu_a \mid a \in \mathbb{X}^n\}$ defined as in corollary 14. Let again $\mathcal{A}_o \subset \mathcal{A}$ denote the algebra of local observables. For any fixed pair (Λ_1, Λ_2) of elements in \mathcal{F} , there exists $a_{12} \in G$ such that $a[\Lambda_1] \bowtie \Lambda_2$ for all $a \in G$ with $|a| > |a_{12}|$. Consequently, by locality (see postulate 11) whenever $a \in G$ with $|a| > |a_{12}|$, $A_1 \in \mathcal{A}_1$ and $A_2 \in \mathcal{A}_2$, we have $\nu_a[A_1]A_2 = A_2\nu_a[A_1]$. By continuity, this entails

COROLLARY 19. *For all $A, B \in \mathcal{A} : \lim_{|a| \rightarrow \infty} \|\nu_a[A]B - B\nu_a[A]\| = 0$, i.e. the group G of translations acts on the algebra \mathcal{A} of quasi-local observables in a norm-asymptotic abelian manner.*

This property makes no sense in the original von Neumann framework for the quantum mechanics of finite systems. In the generalized Haag–Kastler framework devised for infinite systems, this statement which is straightforwardly correct for space translations is rarely satisfied by the time evolution in realistic models that have been controlled.

This raises three questions: the first is whether this property has useful consequences; the second is whether this property can be weakened without jeopardizing the consequences that may be derived from it; and the third is whether any of the weakened forms of this property may be satisfied by the time evolution. I will argue that the answers to the first two questions are “yes”. Specifically, in regard to the first question, see in particular corollary 30 below; and in response to the second, see the forthcoming theorem 25. However, here again, I will warn against the seduction of hypotheses that may ensure a positive answer to the third question, but may be hard to satisfy in specific models; see also the last paragraph in 5.4.B and remark 63(6) below.

DEFINITION 20. A state φ on the algebra \mathcal{A} of quasi-local observables is said to

be translation invariant whenever $\forall (a, A) \in G \times \mathcal{A} : \varphi(\nu_a[A]) = \varphi(A)$, a situation denoted by $\varphi \circ \nu = \varphi$. The state φ is said to be extremal translation invariant if it is translation invariant and may not be written as a convex sum of different translation invariant states.

With G denoting the group of translations of $\mathbb{X}^n = \mathbb{R}^n$ or \mathbb{Z}^n , G is trivially identified with \mathbb{X}^n . Let $\mathcal{C} = \mathcal{C}(G)$ be the set of all complex-valued, continuous, bounded functions $f : G \rightarrow \mathbb{C}$. Henceforth, this set is equipped with the usual point-wise addition and multiplication of functions, and with the sup-norm $\|f\| = \sup_{x \in G} |f(x)|$. These operations equip \mathcal{C} with the structure of an (abelian) C^* -algebra. Define then an action of G on \mathcal{C} by $a[f](x) = f(x - a)$.

DEFINITION 21. With the above notations, an invariant mean on \mathcal{C} is a state η on \mathcal{C} such that $\forall (a, f) \in G \times \mathcal{C} : \eta(a[f]) = \eta(f)$.

Given \mathbb{X}^n , there are several such means. For instance, the ergodic mean on \mathbb{R} may be defined as follows. Let $\mathcal{C}_e = \{f \in \mathcal{C} \mid \lim_{a \rightarrow \infty} 1/2a \int_{-a}^a dx f(x) \text{ exists} \}$. Then $\forall f \in \mathcal{C}_e$, let $\eta_e(f) := \lim_{a \rightarrow \infty} 1/2a \int_{-a}^a dx f(x)$; which then extends by continuity to \mathcal{C} , so as to give an invariant mean, which is the one I will prefer to use in the sequel. One may wish to define similarly the mean η_+ on $\mathcal{C}_+ = \{f \in \mathcal{C} \mid \lim_{x \rightarrow \infty} f(x) \text{ exists} \}$. And, similarly, another mean η_- obtains from the functions that admit a limit as $x \rightarrow -\infty$.

To define averages of states and of observables, notice that for every state φ on the algebra \mathcal{A} of quasi-local observables and any $A, B \in \mathcal{A}$, the functions $\varphi(\nu_\bullet[A]B) : a \in G \mapsto \varphi(\nu_a[A]B) \in \mathbb{C}$ — here the symbol \bullet serves as a reminder to mark the place of the variable a — are continuous and bounded, namely by $\|A\| \|B\|$. Thus the functions $\varphi(\nu_\bullet[A]B)$ belong to \mathcal{C} . When $B = I$ we write simply $\varphi(\nu_\bullet[A])$ for $\varphi(\nu_\bullet[A]I)$. With these notations, the following definition makes sense.

DEFINITION 22. Given an invariant mean η on \mathcal{C} and any state φ on the algebra \mathcal{A} of quasi-local observables, the average $\eta[\varphi]$ of the state φ is defined as the translation invariant state

$$\eta[\varphi] : A \in \mathcal{A} \mapsto \eta(\varphi(\nu_\bullet[A])) \in \mathbb{C} \quad .$$

A translation invariant state φ is said to be η -clustering whenever

$$\forall A, B \in \mathcal{A} : \eta(\varphi(\nu_\bullet[A] B)) = \varphi(A) \varphi(B) \quad .$$

Warnings concerning terminology:

1. η -clustering is also referred to as “weak clustering”.
2. η -clustering should not be confused with the stronger property called “weak mixing”, namely

$$\forall A, B \in \mathcal{A} : \eta |\varphi(\nu_\bullet[A] B) - \varphi(A) \varphi(B)| = 0$$

where for any complex number z , $|z|$ denotes absolute value of z . The name “weak mixing” conforms to the usage in classical ergodic theory, cf. e.g. [Arnold and Avez, 1968, p. 21].

3. The property simply called *clustering* does not involve averaging, and thus is stronger; it is:

$$\forall a \in \mathbb{R}^n \text{ and } \forall A, B \in \mathcal{A} : \lim_{\lambda \rightarrow \infty} \varphi(\nu_{\lambda a}[A] B) = \varphi(A) \varphi(B) \quad .$$

This property is called “mixing” in classical ergodic theory, cf. e.g. [Arnold and Avez, 1968, p. 20].

4. An even stronger property is introduced in definition 27 below.
5. Each of the above properties expresses how much the correlations between $\nu_a[A]$ and B decay with large distances $|a|$ when the system is in the state φ . The term “clustering” affixed to these properties, also used in QFT, seems to be inherited from scattering theory where it expresses the asymptotic independence of separate scattering products, or “clusters”.

The definition of the average of an observable is a little bit more involved. For the general mathematical framework, cf. e.g. [Emch, 1972a, subsection 2.2.d]; in particular, for the general statements and proofs corresponding to scholium 23 and theorem 25 below, cf. [Emch, 1972a, lemma, pp. 174–175] and [Emch, 1972a, theorem 8, pp. 183–184]. Note that, here, the asymptotic abelianness of the action of the group of space translations — corollary 19 above — allows the simpler presentation offered below. *This is where global observables* — cf. definition 15 above — *enter the picture*.

Let φ be a translation invariant state on the algebra \mathcal{A} of quasi-local observables, and $\{\pi_\varphi, \mathcal{H}, \Phi\}$ be the GNS triple associated to φ . Let further $\mathcal{N}_\varphi = \pi_\varphi(\mathcal{A})''$ and $\mathcal{Z}_\varphi = \mathcal{N}_\varphi \cap \mathcal{N}'_\varphi$.

For $a \in G$ fixed, and A running over \mathcal{A} , the map $\pi_\varphi(A)\Phi \in \mathcal{H} \mapsto \pi_\varphi(\nu_a[A])\Phi \in \mathcal{H}$ extends uniquely to a unitary operator $U_a \in \mathcal{U}(\mathcal{H}) := \{U \in \mathcal{B}(\mathcal{H}) \mid U^*U = UU^* = I\}$. This defines a continuous unitary representation $U : a \in G \mapsto U_a \in \mathcal{U}(\mathcal{H})$ such that $\forall (a, A) \in G \times \mathcal{A} : U_a \pi_\varphi(A) U_a^* = \pi_\varphi(\nu_a[A])$.

As usual, let $U(G)' := \{B \in \mathcal{B}(\mathcal{H}) \mid \forall a \in G : U_a B = B U_a\}$ denote the commutant of $U(G)$. Equivalently here, $U(G)' = \{B \in \mathcal{B}(\mathcal{H}) \mid \forall a \in G : U_a B U_a^* = B\}$.

Finally, let $\mathcal{P} := \{\Psi \in \mathcal{H} \mid \forall a \in G : U_a \Psi = \Psi\}$; and denote by P the orthogonal projector from \mathcal{H} onto \mathcal{P} .

SCHOLIUM 23. For every invariant mean η on \mathcal{C} , the map

$$\eta_\varphi : A \in \mathcal{A} \mapsto \eta_\varphi[A] \in \mathcal{Z}_\varphi \cap U(G)'$$

defined, for all $A \in \mathcal{A}$ by

$$\forall \Psi_1, \Psi_2 \in \mathcal{H} : (\Psi_1, \eta_\varphi[A]\Psi_2) = \eta(\Psi_1, \pi_\varphi(\nu_\bullet[A])\Psi_2)$$

is a $*$ -homomorphism and satisfies $\eta_\varphi[A]P = P\eta_\varphi[A] = P\eta_\varphi[A]P$.

DEFINITION 24. Let η be an invariant mean on \mathcal{C} ; φ be a translation invariant state on the algebra \mathcal{A} of quasi-local observables; $\mathcal{N}_\varphi = \pi_\varphi(\mathcal{A})''$ be the algebra of global observables associated to the state φ , via the GNS triple $\{\pi_\varphi, \mathcal{H}, \Phi\}$; and $\mathcal{N}_\varphi^G = \{N \in \mathcal{N}_\varphi \mid \forall a \in G : U_a N U_a^* = N\}$ be the algebra of translation invariant global observables. Then the average of a quasi-local observable $A \in \mathcal{A}$ is defined as the translation invariant global observable $\eta_\varphi[A] \in \mathcal{N}_\varphi^G$.

We are now ready to enunciate the central quantum ergodic theorem relative to the action of the group of space translations.

THEOREM 25. Let $\nu : a \in G \rightarrow \text{Aut}(\mathcal{A})$ denote the action of the space-translation group on the algebra \mathcal{A} of quasi-local observables; and let η be any invariant mean on \mathcal{C} . Then the following conditions on a translation invariant state φ on \mathcal{A} are equivalent:

1. φ is extremal translation invariant;
2. φ is η -clustering, i.e. $\forall A, B \in \mathcal{A} : \eta(\varphi(\nu_\bullet[A]B)) = \varphi(A)\varphi(B)$;
3. the canonical extension $\tilde{\varphi} : N \in \mathcal{N}_\varphi \mapsto (\Phi, N\Phi) \in \mathbb{C}$ of φ to the von Neumann algebra \mathcal{N}_φ of global observables associated to φ is the only translation invariant normal state on this algebra;
4. the invariant subspace $\mathcal{P} \subset \mathcal{H}$ is one-dimensional;
5. the average $\eta_\varphi[A]$ of every quasi-local observable $A \in \mathcal{A}$ is a multiple of the identity, namely $\eta_\varphi[A] = \varphi(A)I$;
6. all translation invariant global observables $N \in \mathcal{N}_\varphi^G := \mathcal{N}_\varphi \cap U(G)'$ are multiples of the identity;
7. $\mathcal{Z}_\varphi \cap U(G)' = \mathbb{C}I$ where $\mathcal{Z}_\varphi := \mathcal{N}_\varphi \cap \mathcal{N}_\varphi'$.

REMARKS 26.

1. Recall that in definition 15 three kinds of observables were introduced. The *local* observables relative to some finite region Λ are described in the original von Neumann formalism [von Neumann, 1932c] where, typically, $\mathcal{A}_\Lambda = \mathcal{B}(\mathcal{H}_\Lambda)$, and $\mathcal{H}_\Lambda = \mathcal{L}^2(\Lambda, dx)$. Thus one refers to local observables as self-adjoint elements of $\mathcal{A}_o = \cup_{\Lambda \in \mathcal{F}} \mathcal{A}_\Lambda$. The *quasi-local observables*, defined abstractly as observables that are norm-limits of local observables, pertain to the microscopic description of many-body systems that are infinitely extended in space; section 5 below opens with three concrete QSP examples. These ‘quasi-local’ observables belong to the C^* -algebra $\mathcal{A} = \overline{\mathcal{A}_o}^n$. Observables of the third kind, the *global observables*, appear at the macroscopic level when bulk properties of matter are investigated; they belong to the von Neumann algebra $\mathcal{N}_\varphi := \pi_\varphi(\mathcal{A})''$ obtained as the weak-closure of the

GNS representation π_φ (of \mathcal{A}) corresponding to a state φ (on \mathcal{A}) specifically obtained by a process called the thermodynamical limit, several examples of which are discussed in the following sections.

Space averages are examples of such *global observables*. A concrete example in ferromagnetism obtains with any one of the three components of the magnetization. Observables of this third kind depend on the global state of the system considered, thus reflecting the preparation of the system. For instance, when the state is extremal translation invariant, these observables are multiples of the identity operator — recall the equivalence of conditions (1) and (5) in theorem 25 — hence their value is the same in all configurations that differ only locally from the given state. Their assuming different values in configurations that differ globally from one another serves as witness for the existence of different thermodynamical phases; cf. subsection 5.7.

2. A global state φ on $\mathcal{A} = \overline{\mathcal{A}_o}$ with $\mathcal{A}_o = \bigcup_{\Lambda \in \mathcal{F}} \mathcal{A}_\Lambda$ is usually defined by continuity from

$$(59) \quad \forall \Lambda \text{ and } \forall A_\Lambda \in \mathcal{A}_\Lambda \quad : \quad \varphi(A_\Lambda) = \lim_{\substack{|\Omega| \rightarrow \infty, \\ \Omega \in \mathcal{F}, \Omega \supseteq \Lambda}} \varphi_\Omega(A_\Lambda)$$

where $\{\varphi_\Omega | \Omega \in \mathcal{F}\}$ is a consistent family of local states. The local states are themselves defined with respect to some consistent boundary conditions; e.g. periodic boundary conditions on every Λ . Hence, the global state φ and thus the von Neumann algebra $\mathcal{N}_\varphi := \pi_\varphi(\mathcal{A})''$ of global observables may depend on the boundary conditions one has chosen. This happens in particular in the presence of the long-range order that often accompanies the onset of phase transitions. This dependence on initial conditions, *even in the thermodynamical limit*, is an ubiquitous phenomenon, known already in classical statistical physics.

Indeed, in an argument that was later confirmed to be correct — for references, cf. e.g. [Emch and Liu, 2002, pp. 416–417] — Peierls [1936] pointed out the fact that the Ising model in two dimensions develops, for sufficiently low temperatures, a sensitivity to boundary conditions: one phase — say the one with strictly positive magnetization — may be selected by clamping all spins on the boundary in the “up” position.

3. Here again, in the special case where $\{\mathcal{N}, \varphi\}$ is $\{\mathcal{L}^\infty(\Omega), \mu\}$, the above theorem reduces to the known classical case. Note however that the theorem is stated here for *space translations* rather than for the *time evolution*; the reason is that the proof uses asymptotic abelianness which space translations satisfy — see corollary 19 above — or some weakened form such as (61) in remark 31 below. Yet, even such a weakened form of asymptotic abelianness is hard to come by for the time evolution of quantum dynamical models.

The clustering condition (2) in the theorem may be strengthened when the representation π_φ is primary, i.e. when the center $\mathcal{Z}_\varphi := \pi_\varphi(\mathcal{A})'' \cap \pi_\varphi(\mathcal{A})'$ satisfies $\mathcal{Z}_\varphi = \mathbb{C}I$. Specifically, for any region $\Lambda \in \mathcal{F}$, let

$$\mathcal{A}_\Lambda^c := \overline{\cup_{\Omega \in \mathcal{F}; \Omega \not\propto \Lambda} \mathcal{A}_\Omega},$$

where, for any subset $\mathcal{B} \subset \mathcal{A}$, $\overline{\mathcal{B}}$ denotes the closure of \mathcal{B} in the norm-topology of \mathcal{A} . As a consequence of locality $A \in \mathcal{A}_\Lambda$ and $B \in \mathcal{A}_\Lambda^c$ entail $AB - BA = 0$. Let now $\mathcal{N}_{\varphi, \Lambda}^c := \pi_\varphi(\mathcal{A}_\Lambda^c)''$.

DEFINITION 27. A state φ on the algebra \mathcal{A} of quasi-local observables is said to be uniformly clustering whenever for any $A \in \mathcal{A}$ and every $\epsilon > 0$, there exists a region $\Lambda \in \mathcal{F}$ depending on A and ϵ , such that

$$(60) \quad \forall B \in \mathcal{A}_\Lambda^c : |\varphi(AB) - \varphi(A)\varphi(B)| \leq \epsilon \|B\| \quad .$$

DEFINITION 28. The elements of the von Neumann algebra $\mathcal{N}_\varphi^\infty := \bigcap_{\Lambda \in \mathcal{F}} \mathcal{N}_{\varphi, \Lambda}^c$ are called observables at infinity with respect to φ .

SCHOLIUM 29. For each state φ separately, the observables at infinity are central, i.e. $\mathcal{N}_\varphi^\infty \subseteq \mathcal{Z}_\varphi$. Moreover the following two conditions on a state φ are equivalent:

1. all observables at infinity are multiples of the identity operator, i.e. $\mathcal{N}_\varphi^\infty = \mathbb{C}I$;
2. φ is uniformly clustering.

Note that definitions 27, 28 and scholium 29 do not require that φ be space-translation invariant, although they involve in an essential manner the local structure of \mathcal{A} . For space-translation invariant states one has in addition:

COROLLARY 30. *The following two conditions:*

1. φ is a translation invariant state on the algebra \mathcal{A} of quasi-local observables;
2. the algebra \mathcal{N}_φ of global observables is a factor, i.e. $\mathcal{Z}_\varphi = \mathbb{C}I$

jointly entail that

- a. φ is extremal translation invariant (and so satisfies the equivalent conditions noted in theorem 25);
- b. φ is uniformly clustering.

REMARKS 31.

1. Condition (2) in corollary 30 is satisfied whenever φ is an extremal KMS state; cf. subsection 5.6 below.
2. The proofs of theorem 25, scholium 29, and corollary 30 are not trivial, but they were all known by the early 1970s; cf. e.g. [Emch, 1972a, theorem II.2.8 and theorem IV.1.7].

3. In particular, the proof of theorem 25 shows that the equivalence of its seven conditions may be obtained in more general contexts where the action of the group of space translations is replaced by an action with respect to which the invariant state φ satisfies the condition of η -abelianness, namely the condition:

$$(61) \quad \forall A, B, C \in \mathcal{A} : \eta \{ \varphi(C^* [\nu_g[A] B - B \nu_g[A]] C) \} = 0 \quad .$$

This condition is much weaker than the norm-asymptotic abelianness proven in corollary 19 for the action of the translation group.

4. It is therefore tempting to try and transfer the above considerations to the group \mathbb{R} governing the time evolution of a quantum dynamical system. In fact if φ is an extremal \mathbb{R} -invariant state, then such a dynamical system will be η -abelian in the sense of (61), *provided* the vector Φ of the GNS representation — which, by construction, is cyclic for $\pi_\varphi(\mathcal{A})$ — is also cyclic for the von Neumann algebra $\pi_\varphi(\mathcal{A})'$, a condition equivalent to the requirement that Φ be separating for the von Neumann algebra $\mathcal{N}_\varphi := \pi_\varphi(\mathcal{A})''$, i.e. $N \in \mathcal{N}_\varphi$ and $N\Phi = 0$ entail $N = 0$. The condition that a von Neumann algebra \mathcal{N} admits a vector Φ that is cyclic for both \mathcal{N} and \mathcal{N}' is referred to by saying that this von Neumann algebra is in standard form; for the relevance of this condition in the present context cf. definition 36 and theorem 39 below. This however only raises again the question of whether φ is extremal under the evolution responsible for the approach to equilibrium. In this respect, we may note that this is the case for the dilated evolution in the example of a chain of weakly coupled harmonic oscillators, discussed at the beginning of this subsection, and in general for the evolution α of generalized Kolmogorov flows; cf. definition 17; see nevertheless the “warning” following remark 18, or subsection 5.4(B).
5. On the mathematical side, quantum ergodic theory may be concerned with group actions more general than space or time translations. In fact, theorem 25 and the third remark just above extend without modifications to the actions of *amenable* groups, i.e. groups \mathcal{G} that admit an invariant mean in the sense of definition 21 (where $G = \mathbb{R}^n$ or \mathbb{Z}^n is replaced by \mathcal{G}). For a general presentation of the theory of amenable groups, cf. e.g. [Greenleaf, 1969] or for a brief review geared to applications in QSP [Emch, 1972a, pp. 164–172]. Restricting attention here to locally compact groups, let it suffice to note that compact groups, abelian groups, and semi-direct products thereof are amenable; in particular the rotation groups, translation groups, and Euclidean groups in finite-dimensional Euclidean spaces are amenable. However, *no* non-compact semi-simple Lie group is amenable, so that in particular the Lorentz group of 4-dimensional relativistic QFT is *not* amenable.
6. Pushing the theory even further than amenable group actions may be done by considering “large groups of automorphisms” of a C^* -algebra \mathcal{A} , i.e.

actions $\alpha : \mathcal{G} \rightarrow \text{Aut}(\mathcal{A})$ that satisfy for every self-adjoint $A \in \mathcal{A}$ and every \mathcal{G} -invariant state φ on \mathcal{A} :

$$(62) \quad {}^{w-op}\overline{\text{co}\{\pi_\varphi(\alpha_g[A]) \mid g \in \mathcal{G}\}} \cap \pi_\varphi(\mathcal{A})' \neq \emptyset \quad ,$$

where for any subset S of a vector space, $\text{co}\{S\}$ denotes the ‘‘convex hull’’ of S , i.e. the collection of all convex combinations of elements in S ; and for any set $\mathcal{B} \subset \mathcal{B}(\mathcal{H})$, ${}^{w-op}\overline{\mathcal{B}}$ denotes the closure of \mathcal{B} in the weak-operator topology of $\mathcal{B}(\mathcal{H})$. The notion of *large group of automorphisms* was introduced by St ormer in 1967 who used it soon afterwards to prove a quantum analogue of de Finetti’s exchangeability theorem in classical probability theory [St ormer, 1969]; for a review and some applications to the semantic foundations of quantum theory, cf. e.g. [Emch, 2005] and references therein. Note that any amenable group action for which the system is η -abelian for some mean η is a large group of automorphisms for this system.

Here again, one can hardly resist the conclusion that quantum ergodic theory is now a mature mathematical theory in search of further physical applications to QSP, most notably through the understanding it provides for the various clustering (or mixing) properties described in the present section; cf. e.g. subsections 5.4 and 5.7 below.

4 THE KMS CONDITION FOR EQUILIBRIUM

The identification of the KMS condition as a canonical characterization of equilibrium states appears in the confluence of two currents of thought.

The first source is the recognition by Kubo [1957] and by Martin & Schwinger [1959] that objects which play a central role in condensed matter physics, namely the so-called thermal Green functions — cf. e.g. [Bonch-Bruевич and Tyablikov, 1962] — possess remarkable analytic properties. For a foretaste, see scholium 32 below.

The second source of inspiration is recognizable in the original texts [Murray and von Neumann, 1936] of what was to become the theory of von Neumann algebras, and is emphasized in the candid reminiscences of one of the pioneers of this theory [Murray, 1990]. A great deal of the theory could be built from the following observation: there are matrix algebras \mathcal{N} which, together with their commutant \mathcal{N}' , satisfy the following properties:

(i) they are factors, i.e. have trivial center: $\mathcal{N} \cap \mathcal{N}' = \mathbb{C}I$; (ii) \mathcal{N} and \mathcal{N}' admit a common cyclic vector Φ ; (iii) there exists an involutive antiunitary operator J such that $J\Phi = \Phi$ and $N \in \mathcal{N} \mapsto JNJ \in \mathcal{N}'$ is bijective. For a concrete, simple example, see equation (71) below.

Each of the two facets of the theory — analytic and algebraic — involves some mathematical intricacies; hence the division of this section into two subsections: first, a simple example; and second, the general theory.

4.1 A Wignerian Approach

In this subsection, I wish to abide by Wigner’s famous dictum [Wigner, 1962]: “Please explain it with 2×2 matrices.” Accordingly I proceed with the description of what happens to a quantum $1/2$ -spin in canonical equilibrium at natural temperature $\beta > 0$ in a magnetic field B parallel to the z -axis. The observables are the self-adjoint elements of the algebra \mathcal{M} of 2×2 matrices with complex entries. The Hamiltonian is

$$(63) \quad H = -B\sigma^z = \begin{pmatrix} \epsilon_1 & 0 \\ 0 & \epsilon_2 \end{pmatrix} \quad \text{with} \quad \epsilon_1 = -B, \quad \epsilon_2 = +B.$$

The canonical equilibrium state is, according to von Neumann’s characterization (38):

$$(64) \quad \psi_H : M \in \mathcal{M} \rightarrow \text{Tr}(\rho_H M) \quad \text{with} \quad \rho_H = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

where $\lambda_n = Z_H^{-1} \exp(-\beta\epsilon_n)$, with $Z_H = \exp(-\beta\epsilon_1) + \exp(-\beta\epsilon_2)$ denoting the partition function of the system.

In the Heisenberg picture, conjugate to Schrödinger’s picture (39), the evolution is

$$(65) \quad \left. \begin{aligned} \alpha_t : M \in \mathcal{M} \rightarrow \alpha_t[M] = U^*(t)MU^*(-t) \in \mathcal{M} \\ \text{with} \quad U^*(t) = \begin{pmatrix} e^{i\epsilon_1 t} & 0 \\ 0 & e^{i\beta\epsilon_2 t} \end{pmatrix} . \end{aligned} \right\}$$

To make computations easier and, moreover, immediately generalizable to higher dimensions, consider the matrices

$$E_{mn} : \Psi \in \mathbb{C}^2 \mapsto (\Psi_n, \Psi)\Psi_m \in \mathbb{C}^2$$

where $\{\Psi_n \mid n = 1, 2\}$ are eigenvectors of H , i.e. with

$$\Psi_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} ; \quad \Psi_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix} :$$

$$E_{11} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} ; \quad E_{12} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} ; \quad E_{21} = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} ; \quad E_{22} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} .$$

These matrices form a basis in \mathcal{M} and — with ψ_H and α_t as in (64) and (65) — satisfy

$$E_{kl}E_{mn} = \delta_{lm}E_{kn} \quad , \quad \psi_H(E_{mn}) = \lambda_m\delta_{mn} \quad , \quad \alpha_t(E_{mn}) = e^{i(\epsilon_m - \epsilon_n)t}E_{mn} .$$

From these relations and the identity $\exp[-\beta(\epsilon_m - \epsilon_n)]\lambda_n = \lambda_m$, one obtains that the analytic functions

$$f_{klmn} : z \in \mathbb{C} \rightarrow \lambda_n e^{i(\epsilon_m - \epsilon_n)z} \delta_{lm} \delta_{kn}$$

satisfy $\forall t \in \mathbb{R} : f_{klmn}(t) = \psi_H(E_{kl}\alpha_t[E_{mn}])$ and $f_{klmn}(t+i\beta) = \psi_H(\alpha_t[E_{mn}]E_{kl})$. Moreover, on the strip

$$\Omega_\beta := \{z \in \mathbb{C} \mid 0 \leq \text{Im } z \leq \beta\},$$

the analytic functions f_{klmn} are bounded, namely by $\exp(|\epsilon_m - \epsilon_n|\beta)$.

These two properties of the canonical equilibrium state ψ_H extend by linearity to the time correlation functions

$$(66) \quad f_{MN}(t) = \psi_H(M\alpha_t[N]) \quad \text{and} \quad f_{MN}(t+i\beta) = \psi_H(\alpha_t[N]M)$$

with M and N arbitrary in \mathcal{M} .

Conversely, suppose that φ is a state on \mathcal{M} such that for every pair M, N of elements in \mathcal{M} there exists a function $f_{M,N} : z \in \Omega_\beta \mapsto f_{M,N}(z) \in \mathbb{C}$ such that

- (i) $f_{M,N}$ is bounded and continuous on the strip Ω_β ;
- (ii) $f_{M,N}$ is analytic inside that strip;
- (iii) for all $t \in \mathbb{R} : f_{M,N}(t) = \varphi(M\alpha_t[N])$ and $f_{M,N}(t+i\beta) = \varphi(\alpha_t[N]M)$.

Then in particular, with $M = I$, the function $f_{I,N}$ is periodic with period $i\beta$. It may then be extended to a function that is both bounded and analytic on the whole complex plane. The classical Liouville's theorem — cf. e.g. [Churchill and Brown, 1990, theorem 43.1] — thus entails that this function must be constant, i.e. for all $(t, N) \in \mathbb{R} \times \mathcal{M} : \varphi(\alpha_t[N]) := \text{Tr } U^*(-t)\rho U^*(t)N$ is equal to $\text{Tr } \rho N = \varphi(N)$; and thus

$$\rho = \begin{pmatrix} \mu_1 & 0 \\ 0 & \mu_2 \end{pmatrix}$$

where the values of μ_1, μ_2 positive with $\mu_1 + \mu_2 = 1$ are computed presently. Upon comparing, for every pair of indices (m, n) the analytic continuation of

$$f_{mn}(t) = \varphi(E_{nm}\alpha_t[E_{mn}]) = e^{i(\epsilon_m - \epsilon_n)t} \mu_n$$

and

$$f_{mn}(t+i\beta) = \varphi(\alpha_t[E_{mn}]E_{nm}) = e^{i(\epsilon_m - \epsilon_n)t} \mu_m$$

one obtains $\exp[-\beta(\epsilon_m - \epsilon_n)]\mu_m = \mu_n$ and thus, upon imposing the normalization $\varphi(I) = 1$, i.e. $\mu_1 + \mu_2 = 1$:

$$\mu_n = \frac{e^{-\beta\mu_n}}{e^{-\beta\mu_1} + e^{-\beta\mu_2}} = \lambda_n \quad .$$

Thus, indeed $\varphi = \psi_H$.

In summary, one obtained by elementary means an elementary illustration of the first facet of the theory, its analytic aspect:

SCHOLIUM 32. Let $H = -B\sigma^z$ be the Hamiltonian describing a spin $\frac{1}{2}$ in a magnetic field B . Then, for any state φ on $\mathcal{M} = \mathcal{M}(2, \mathbb{C})$, the following conditions are equivalent:

- (I) φ is the canonical equilibrium state ψ_H with respect to the Hamiltonian H ;
- (II) for every pair (M, N) of elements of \mathcal{M} there exists a function $f_{M,N} : z \in \Omega_\beta \rightarrow \mathbb{C}$ such that

$$(67) \quad \left. \begin{array}{l} f_{M,N} \text{ is bounded and continuous on } \Omega_\beta ; \\ f_{M,N} \text{ is analytic in the interior of } \Omega_\beta ; \\ \forall t \in \mathbb{R} : \begin{cases} f_{M,N}(t) &= \varphi(M \alpha_t[N]) \\ f_{M,N}(t + i\beta) &= \varphi(\alpha_t[N] M) \end{cases} \end{array} \right\} .$$

Moving now towards the algebraic aspect of the theory, one pursues with the same simple model, and let φ be a faithful state over \mathcal{M} , i.e. a state such that $M \in \mathcal{M}$ and $\varphi(M^*M) = 0$ entail $M = 0$. Without loss of generality one may choose a basis in which the density matrix ρ corresponding to φ is diagonal, with eigenvalues λ_n ($n = 1, 2$) strictly positive since φ is supposed to be faithful. Consider the representation π of \mathcal{M} given by:

$$(68) \quad \forall M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathcal{M} \quad : \quad \pi(M) = \begin{pmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & c & d \end{pmatrix} = M \otimes I$$

acting on the Hilbert space \mathbb{C}^4 equipped with its standard scalar product in which Ψ_{kl} defined by

$$\Psi_{11} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad \Psi_{21} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad \Psi_{12} = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \quad \Psi_{22} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

is an orthonormal basis. The vector

$$(69) \quad \Phi = \sum_k \lambda_k^{\frac{1}{2}} \Psi_{kk} = \begin{pmatrix} \lambda_1^{\frac{1}{2}} \\ 0 \\ 0 \\ \lambda_2^{\frac{1}{2}} \end{pmatrix}$$

satisfies $\Psi_{mn} = \lambda_n^{\frac{1}{2}} \pi(E_{mn})\Phi$, from which one reads:

$$\mathbb{C}^4 = \{ \pi(M)\Phi \mid M \in \mathcal{M} \} \quad \text{and} \quad \forall M \in \mathcal{M} : (\Phi, \pi(M)\Phi) = \varphi(M) \quad .$$

Hence $\{\mathcal{H} := \mathbb{C}^4, \pi, \Phi\}$ is the canonical GNS triple associated to the state φ . Moreover, since φ is assumed to be faithful, $\|\pi(M)\Phi\| = 0$ entails $M = 0$, i.e. Φ is also separating for $\pi(\mathcal{M})$. The essential step now is to introduce the two operators J and Δ defined on \mathcal{H} by the conditions that J is antilinear, Δ is linear, with

$$J\Psi_{mn} = \Psi_{nm} \quad \text{and} \quad \Delta\Psi_{mn} = \frac{\lambda_m}{\lambda_n} \Psi_{mn} \quad .$$

Note that, since Δ is given here with its spectral resolution, the functions of this operator may be defined by linearity from $f(\Delta) : \Psi_{mn} \in \mathbb{C}^4 \mapsto f(\frac{\lambda_m}{\lambda_n})\Psi_{mn} \in \mathbb{C}^4$. In particular, $\{\Delta^{is} | s \in \mathbb{R}\}$ is a continuous group of unitary operators acting on \mathbb{C}^4 .

One verifies immediately from their definition above that the operators J and Δ satisfy the following properties. Firstly,

$$(70) \quad J \text{ is an isometry, } J^2 = I, \Delta \text{ is self-adjoint, } J\Delta J = \Delta^{-1}, J\Phi = \Phi = \Delta\Phi.$$

Secondly,

$$(71) \quad J \begin{pmatrix} a & b & 0 & 0 \\ c & d & 0 & 0 \\ 0 & 0 & a & b \\ 0 & 0 & c & d \end{pmatrix} J = \begin{pmatrix} a^* & 0 & b^* & 0 \\ 0 & a^* & 0 & b^* \\ c^* & 0 & d^* & 0 \\ 0 & c^* & 0 & d^* \end{pmatrix} \in I \otimes \mathcal{M}$$

i.e. upon denoting by \mathcal{N} the image $\pi(\mathcal{M})$ of \mathcal{M} through the representation π , we have: $J\mathcal{N}J = \mathcal{N}'$; hence (71) gives an explicit bijection from \mathcal{N} onto its commutant \mathcal{N}' . The relation (71) is a particular case of the general Tomita-Takesaki duality (see theorem 39 below).

Thirdly, with $\beta > 0$ arbitrary, but fixed, we have $\forall t \in \mathbb{R} : \Delta^{-it/\beta} \Psi_{mn} = \exp[i(\epsilon_m - \epsilon_n)t] \Psi_{mn}$. Hence $\Delta^{-it/\beta} \pi(E_{mn}) \Delta^{it/\beta} \Psi_{kl} = \exp[i(\epsilon_m - \epsilon_n)t] \pi(E_{mn}) \Psi_{kl}$ with $\epsilon_n = c - (1/\beta) \ln \lambda_n$ where c is an arbitrary real constant. Consequently, the unitary group $\{\Delta^{it/\beta} | t \in \mathbb{R}\}$ implements a group of automorphisms of \mathcal{N} , namely

$$(72) \quad \tau_t : N \in \mathcal{N} \rightarrow \tau_t[N] = \Delta^{-it/\beta} N \Delta^{it/\beta} \in \mathcal{N}$$

with, for all $(t, M) \in \mathbb{R} \times \mathcal{M}$, $\tau_t[\pi(M)] = \pi(\alpha_t[M])$ with $\alpha_t[M] = \exp^{iHt} M \exp^{-iHt}$ and $H = \sum_n \epsilon_n E_{nn}$. Summing up, this establishes that φ is the canonical equilibrium state at natural temperature β for the Hamiltonian H just constructed.

Fourthly, the operator $S = J\Delta^{\frac{1}{2}}$ satisfies $S\pi(E_{mn})\Phi = \pi(E_{nm})\Phi$ and thus, since J and therefore S are antilinear:

$$(73) \quad \forall N \in \mathcal{N} : \quad S N \Phi = N^* \Phi \quad .$$

Finally, the generator L of the unitary group $\{\Delta^{it/\beta} | t \in T\}$ on $\mathbb{C}^4 = \mathbb{C}^2 \otimes \mathbb{C}^2$ is

$$(74) \quad L = H \otimes I - I \otimes H,$$

so that the spectrum of L is symmetric around 0 : $Sp(L) = \{\epsilon_2 - \epsilon_1, 0, 0, \epsilon_1 - \epsilon_2\}$.

SCHOLIUM 33. Let $\{\mathbb{C}^4, \pi, \Phi\}$ be the GNS triple canonically associated to a faithful state φ on the algebra \mathcal{M} of 2×2 matrices; and let \mathcal{N} be the von Neumann algebra $\pi(\mathcal{M}) = \{\pi(M) | M \in \mathcal{M}\}$ acting on $\mathcal{H} = \mathbb{C}^4$. Then

1. \mathcal{N} is isomorphic to \mathcal{M} and φ may be regarded as a faithful state on \mathcal{N} ;

2. Φ is both cyclic and separating for \mathcal{N} ;
3. the anti-linear operator defined by $S : N\Phi \in \mathcal{H} = N^*\Phi \in \mathcal{H}$ has polar decomposition $S = J\Delta^{\frac{1}{2}}$ where J is an involutive, anti-linear isometry from \mathcal{H} onto itself, and Δ is a positive operator acting on \mathcal{H} ;
4. J establishes a duality between \mathcal{N} and its commutant; specifically: $N \in \mathcal{N} \mapsto JNJ \in \mathcal{N}'$ is an anti-linear bijection;
5. $\{\Delta^{-it/\beta} \mid t \in \mathbb{R}\}$ implements a group of automorphisms τ_t of \mathcal{N} with respect to which the state φ satisfies the analyticity condition described in scholium 32;
6. $J\Phi = \Phi = \Delta\Phi$, $J^2 = J$ and $\forall s \in \mathbb{R} : J\Delta^{is} = \Delta^{is}J$.

REMARKS 34. Upon surveying the proofs of scholia 32 and 33, one verifies that they can be extended verbatim from $\mathcal{M} = \mathcal{M}(2, \mathbb{C})$ to $\mathcal{M} = \mathcal{M}(n, \mathbb{C})$ where n is any finite positive integer. These scholia extend further to $\mathcal{M} = \mathcal{B}(\mathcal{H})$ where \mathcal{H} is a Hilbert space (with countable basis) provided that:

- (i) the Hamiltonian H satisfies $\text{Tr}(-\beta H) < \infty$;
- (ii) the state φ is countably additive, retaining the condition that φ be faithful.

Indeed, under these circumstances one can read again the proofs of the scholia, now for the *-algebra $\mathcal{E} = \text{Span}\{\pi(E_{mn}) \mid m, n = 1, \dots\}$ of all finite linear combinations of the operators $\pi(E_{mn})$ with $E_{mn} : \Psi \in \mathcal{H} \rightarrow (\Psi_n, \Psi)\Psi_m \in \mathcal{H}$ where again $\{\Psi_n \mid n = 1, 2, \dots\}$ is an orthonormal basis in \mathcal{H} . The extension from \mathcal{E} to the von Neumann algebra $\mathcal{B}(\mathcal{H})$ obtains since the assumptions that φ is countably additive and faithful allow one to use standard continuity arguments, namely here, e.g. [Dixmier, 1957, theorem I.3.5, lemma I.4.4, proposition I.4.1]; or [Kadison and Ringrose, 1983/1986, volume ii, chapter 7]. In particular, $\mathcal{N} = \pi(\mathcal{B}(\mathcal{H})) = \{\pi(\mathcal{M}) \mid M \in \mathcal{B}(\mathcal{H})\}$ is already a von Neumann algebra — i.e. $\mathcal{N} = \mathcal{N}''$ — and is isomorphic to $\mathcal{B}(\mathcal{H})$. Since $\mathcal{B}(\mathcal{H})$ is a factor, so is \mathcal{N} , i.e. the center of this von Neumann algebra is trivial: $\mathcal{N} \cap \mathcal{N}' = \mathbb{C}I$. Moreover \mathcal{N} may be identified with $\mathcal{B}(\mathcal{H}) \otimes \mathbb{C}I$ and \mathcal{N}' with $\mathbb{C}I \otimes \mathcal{B}(\mathcal{H})$.

The von Neumann formalism for quantum mechanics [von Neumann, 1932c] allows one to go this far, but no further. Recall that some of the reasons why one needs to proceed further were indicated in subsection 3.3. The next subsection provides an important tool toward achieving this.

4.2 The Kubo–Martin–Schwinger condition and the Tomita–Takesaki theory

The above results suggest three definitions; the first two are just matters of mathematical terminology, but the third is at the heart of this section.

DEFINITION 35. A state φ on a von Neumann algebra \mathcal{N} is said to be normal whenever it is countably additive., i.e. $\varphi(\sum_n P_n) = \sum_n \varphi(P_n)$ for each countable family $\{P_n\}$ of mutually orthogonal projections in \mathcal{N} .

This simply extends to general von Neumann algebras condition (35), already recognized in [von Neumann, 1932c] as the quantum analogue of the complete additivity of probability measures. The next definition formalizes in the present context some of the notions encountered in the motivating examples covered in the previous subsection.

DEFINITION 36. A von Neumann algebra \mathcal{N} acting on a Hilbert space \mathcal{H} is said to be in standard form whenever there exists a vector $\Phi \in \mathcal{H}$ that is both cyclic and separating for \mathcal{N} , i.e. $\mathcal{N}\Phi$ is norm dense in \mathcal{H} and for $N \in \mathcal{N}$, $N\Phi = 0$ entails $N = 0$.

REMARKS 37. This concept has been around for a long time, but it seems fair to say that full recognition of its central importance in the general theory of von Neumann algebras had to wait for the Tomita–Takesaki modular theory [Tomita, 1967; Takesaki, 1970a]. At the most basic level, notice that if \mathcal{N} is in standard form, one may assume without loss of generality that $\|\Phi\| = 1$, so that $\varphi : N \in \mathcal{N} \rightarrow (\Phi, N\Phi) \in \mathbb{C}$ is a faithful normal state on \mathcal{N} .

Conversely it follows, from the same continuity arguments as those used in remark 34 above, that if φ is any normal state on a von Neumann algebra \mathcal{N} , the GNS representation π corresponding to φ is already a von Neumann algebra; if φ is faithful, then \mathcal{N} is isomorphic to $\pi(\mathcal{N})$. Thus the canonical GNS vector Φ is not only cyclic, but it is also separating. Hence whenever φ is a faithful normal state, \mathcal{N} is isomorphic to $\pi(\mathcal{N})$ which is a von Neumann algebra presented in standard form

The third definition pertains to the core of this section. It is an adaptation of the work of [Kubo, 1957; Martin and Schwinger, 1959], proposed by [Haag *et al.*, 1967] as an extension of the definition of canonical equilibrium states on the global C^* -algebra to be associated to an infinite system.

DEFINITION 38. Let \mathcal{A} be a C^* -algebra, and let $\alpha : t \in \mathbb{R} \rightarrow \alpha_t \in \text{Aut}(\mathcal{A})$ be a group of automorphisms of \mathcal{A} . A state φ on \mathcal{A} is said to satisfy the KMS condition with respect to α for the natural temperature β if for every pair (A, B) of elements of \mathcal{A} there exists a function $f_{A,B}$ defined on the strip $\Omega_\beta = \{z \in \mathbb{C} \mid 0 \leq \text{Im}z \leq \beta\}$, such that $f_{A,B}$ is bounded and continuous on Ω_β ; $f_{A,B}$ is analytic in the interior of Ω_β ; and $\forall t \in \mathbb{R} : f_{A,B}(t) = \varphi(A \alpha_t[B])$ and $f_{A,B}(t + i\beta) = \varphi(\alpha_t[B] A)$.

The main mathematical result of this section, taken from the Tomita–Takesaki modular theory [Tomita, 1967; Takesaki, 1970a], may now be stated.

THEOREM 39 (Tomita–Takesaki). *Let \mathcal{N} be a von Neumann algebra acting on a Hilbert space \mathcal{H} and admitting a cyclic and separating unit vector Φ . Then the closed antilinear operator S obtained as the closure of the map $N\Phi \rightarrow N^*\Phi$, defined for all $N \in \mathcal{N}$, has polar decomposition $S = J\Delta$ where $J = J^2$ is an antilinear isometry from \mathcal{H} onto itself, satisfying $JN'J = \mathcal{N}'$; and Δ is a self-*

adjoint operator (not necessarily bounded!) that is positive, and such that $J\Delta^{it} = \Delta^{it}J$; and for any $\beta > 0 \forall (t, N) \in \mathbb{R} \times \mathcal{N} : \tau_t[N] = \Delta^{-it/\beta} N \Delta^{it/\beta}$ defines a group $\{\tau_t\}$ of $*$ -automorphisms of \mathcal{N} with respect to which the faithful normal state $\varphi : N \in \mathcal{N} \rightarrow (\Phi, N\Phi) \in \mathbb{C}$ satisfies the KMS condition for β . Moreover $\{\tau_t \mid t \in \mathbb{R}\}$ is the unique group of $*$ -automorphisms of \mathcal{N} with respect to which φ satisfies this condition.

REMARKS 40.

1. The theorem generalizes to any arbitrary von Neumann algebra in standard form the result we described — in remark 34 — for the GNS representation of $\mathcal{B}(\mathcal{H})$ associated to any of its faithful normal states.
2. It is essential to the purpose of the present review to emphasize that the theorem does *not* require that \mathcal{N} be a factor.
3. Whereas the theorem asserts that the dynamics τ is uniquely determined by the KMS condition, the converse is *not* true: when \mathcal{N} is not a factor, there exist other normal states on \mathcal{N} that also satisfy the KMS condition with respect to the same dynamics. Indeed, when \mathcal{N} is not a factor, one verifies that for every $Z \neq 0$ that belongs to the center $\mathcal{Z} = \mathcal{N} \cap \mathcal{N}'$, $\psi : N \in \mathcal{N} \rightarrow [\varphi(Z^*Z)]^{-1}\varphi(Z^*NZ)$ defines a normal state that again satisfies the KMS condition with respect to τ for the same β . This remark, the proof of which will be given in subsection 5.6, is essential to the arguments presented in subsection 5.7.
4. Beyond its mathematical attractiveness, the legitimacy of the conjecture that the KMS condition may be regarded as a definition of canonical equilibrium states in the QSP of macroscopic systems will also be discussed in the next section.
5. Finally, mathematical probity requires us to mention that — factor or not — a major difficulty in the proof of theorem 39 resides in showing that the map $N\Phi \rightarrow N^*\Phi$ is closable; for the resolution of this problem, cf. the original papers [Tomita, 1967; Takesaki, 1970a]; it is probably fair to warn the reader that even the didactic presentation in [Kadison and Ringrose, 1983/1986, chapter 9] would have carried us beyond the bounds of this essay. To convey nevertheless an idea of the structures involved in the theorem, I resorted therefore to presenting first the models covered in the preliminary scholia 32 and 33, as these could be treated with mathematically elementary tools. The drawback was however that these models, as well as their routine extensions from $\mathcal{M}(2, \mathbb{C})$ to $\mathcal{B}(\mathcal{H})$ described in remark 34, only involve factors, in fact faithful representations of $\mathcal{B}(\mathcal{H})$, that are not sufficient to cover the macroscopic purposes of QSP where infinitely many degrees of freedom are brought to play. As Haag, Hugenholtz, and Winnink [1967] correctly envisaged, it is the generality involved in theorem 39 that is actually needed in physical applications. The temporal coincidence of this physical intuition

and the arrival on the scene of the mathematical theory of Tomita–Takesaki [1967; 1970a] is a truly remarkable event vividly recounted in [Kadison, 1990, pp. 77–79].

5 KMS CONDITION, QSP AND THERMODYNAMICS

This section presents some of the evidences supporting the physical interpretation of the KMS condition proposed in [Haag *et al.*, 1967] as an alternative definition of equilibrium states in QSP. We already saw that for finite systems the KMS condition is satisfied by the canonical equilibrium states of von Neumann, and only by those states. Now, in subsections 5.1–5.3 models are described to show how the modular structures invented and developed by [Tomita, 1967; Takesaki, 1970a] — which we saw (cf. scholium 33) are realized in finite systems in canonical equilibrium — are also encountered in the equilibrium QSP of infinite systems, thus allowing one to go beyond von Neumann’s formalism [von Neumann, 1932c]. In subsection 5.4 various stability conditions are exhibited that give a thermodynamical characterization of KMS states in QSP. A brief excursion is undertaken in subsection 5.5 to indicate some vistas toward the recognition of the role the KMS condition has later been called to play in relativistic QFT, a role dubbed “revolutionary” by the practitioners. Subsection 5.6 is a mathematical interlude devoted to the algebraic characterization of *extremal* KMS states. When we return to QSP in subsection 5.7, systems that exhibit phase transitions are considered and the unique decomposition of any canonical equilibrium state into its pure thermodynamical phases is shown to be closely modeled by a unique decomposition of KMS states into extremal KMS states. In particular, this subsection is oriented toward substantiating the overarching idea that the KMS condition provides the thermodynamics of infinite systems with a conceptual scheme in which phase transitions occur accompanied by spontaneous symmetry breakdown.

5.1 *Beyond Fock space: The BCS model*

The first indication that something was amiss in the use of the von Neumann formalism in QSP was the Bardeen–Cooper–Schrieffer model for superconductivity, the BCS model. Indeed, in the original treatment of this model [Bardeen *et al.*, 1957], the Hamiltonian chosen to describe a specific interaction between the electrons in a finite but large metallic solid is invariant under gauge transformations of the first kind; an approximation is then proposed, which is asserted to become exact in the infinite volume limit; in this formal process however this symmetry is lost; moreover, the spectrum of the resulting Hamiltonian presents an energy gap that is temperature–dependent. One might argue that the experimentalist may not wish to be concerned with the breaking of that symmetry, but the energy gap cannot be ignored: experimentalists do measure it in the laboratory. Thus, mathematical physicists thought that they ought to understand — how or rather

whether — the Hamiltonian itself may indeed depend on the temperature. Within five years, the culprit was found by Haag [1962] to be that the whole treatment was allegedly carried out in a fixed irreducible representation of the CCR, the then ubiquitous Fock representation, and that this constraint was doing violence to the model.

Specifically, the original Hamiltonian is

$$(75) \quad H_\Lambda = \sum_{p,s} \epsilon(p) a_s(p)^* a_s(p) + \sum_{p,q} b(p)^* \tilde{v}(p,q) b(q)$$

where Λ is the region of space in which the system is contained, typically a cubic box of finite volume $|\Lambda|$; p and q label momentum and are integer multiples of $2\pi|\Lambda|^{-\frac{1}{2}}$; $s = \pm\frac{1}{2}$; $a_s(p)^*$ and $a_s(p)$ are the creation and annihilation operators for an electron of spin s and momentum p ; $\epsilon(p) = -\mu + \frac{1}{2}p^2/2m$ is the energy of a free electron of momentum p ; $b(p)^* = a_\uparrow(p)^* a_\downarrow(-p)^*$ is the creation operator of a so-called Cooper pair; and $b(p)^* \tilde{v}(p,q) b(q)$ is the interaction energy between two Cooper pairs, i.e. four electrons, so that the Hamiltonian (75) is *not* quadratic in the original field operators. The form of $\tilde{v}(p,q)$ will be discussed later on.

The approximating Hamiltonian is

$$(76) \quad \tilde{H}_\Lambda = \sum_{p,s} E(p) \gamma_s(p)^* \gamma_s(p)$$

where $\gamma_s(p)^*$ and $\gamma_s(p)$ are the creation and annihilation operators for the elementary excitations given by a Bogoliubov–Valatin transformation

$$(77) \quad \left. \begin{aligned} \gamma_\uparrow(p) &= u(p) a_\uparrow(p) + v(p) a_\downarrow(-p)^* \\ \gamma_\downarrow(p) &= -v(-p) a_\uparrow(-p)^* + u(-p) a_\downarrow(p) \end{aligned} \right\}$$

where

$$(78) \quad \begin{aligned} E(p) &= \{\epsilon(p)^2 + [\Delta(p)\Delta(p)^*]\}^{\frac{1}{2}} & u(p) &= \Delta(p)^*/D(p) \\ D(p) &= \{[E(p) - \epsilon]^2 + [\Delta(p)\Delta(p)^*]\}^{\frac{1}{2}} & v(p) &= [E(p) - \epsilon(p)]/D(p) \end{aligned} ;$$

and Δ satisfies the all-important self-consistency equation

$$(79) \quad \Delta(p) = - \sum_q \tilde{v}(p,q) \frac{\Delta(q)}{2E(q)} \tanh\left(\frac{1}{2}\beta E(q)\right) .$$

Clearly $\Delta = 0$ is always a solution, in which case the spectra of H and \tilde{H} coincide; this is the normal phase in which nothing particularly interesting happens. The essence of the model is that there is a critical temperature T_c (recall $\beta = 1/kT$) below which an energetically more favorable solution $\Delta \neq 0$ develops. This corresponds to the superconducting phase. We henceforth pursue the discussion for $0 < T < T_c$.

This is the phase we are interested in, and it may be useful to recall in physical terms what the physicists first saw in (76)–(79). BCS devised a limiting procedure — involving the thermodynamical limit and a “mean-field approximation” (weak,

but very long range interaction) — by which the original Hamiltonian (75) and the new Hamiltonian (76) become interchangeable in the sense that they are claimed to lead to the same limit. While (75) is expressed in terms of the electrons' creation and annihilation operators $a_s^b(p)$; the new Hamiltonian (76) is free in terms of the elementary excitations $\gamma_s^b(p)$. The energy spectrum of these excitations is $\{E(p)\}$ and differs — see (78) — from the energy spectrum $\{\epsilon(p)\}$ of the free electrons by a temperature-dependent “gap” which is observable in the laboratory; the numerical results so obtained for this gap are in very good agreement with the prediction (79); cf. [Schrieffer, 1974, Figure 1–3].

The mathematical picture however demands some explanation. Indeed: (i) the initial Hamiltonian (75) is invariant under the gauge symmetry defined, for any $\theta \in (0, 2\pi]$ by $a_s(p) \rightarrow \exp(i\theta)a_s(p)$ whereas the Hamiltonian (76) is not; and (ii) the energy spectrum $\{E(p)\}$ of the Hamiltonian (76) is temperature dependent, whereas there is no temperature dependence in (75).

The question therefore is to account for how one could possibly claim — as was done in the prevailing folklore — that such an approximation could become exact in the thermodynamical limit. For this, one has to examine where Δ comes from, namely that $\Delta(p)$ is a scalar multiple of the identity operator, to be viewed as an approximation of the operator $\hat{\Delta}(p) = \sum_q \tilde{v}(p, q)b(q)$. The argument for this is based on the remark that, under suitable assumptions on \tilde{v} , one can arrange for the limit $|\Lambda| \rightarrow \infty$ of $\hat{\Delta}(p)$ to exist — in the weak-operator topology — and to commute with all the creation and annihilation operators $a_s(q)^*$ and $a_s(q)$ which generate an algebra which is tacitly assumed to be irreducible. In this limit, the operator $\Delta(p)$ would be replaced by a scalar multiple of the identity. Some “suitable” assumptions seemed to be achieved when \tilde{v} is the double Fourier transform

$$\tilde{v}(p, q) = \int_{\Lambda} dx dy f(p, x)v(x, y)f(q, y)^* \quad \text{where} \quad f(p, x) = \begin{cases} |\Lambda|^{-\frac{1}{2}} e^{ipx} & x \in \Lambda \\ 0 & x \notin \Lambda \end{cases}$$

with a nonlocal potential v such that $v(x, y)^* = v(y, x)$, $c = \int dx dy |v(x, y)| < \infty$ and $\sum_q |\tilde{v}(p, q)| < \infty$, so that $\lim_{\Lambda \rightarrow \infty} |\tilde{v}(p, q)| = 0$ and $|\tilde{v}(p, q)| \leq c/|\Lambda|$.

The practitioner will recognize here an approximation of the mean molecular field type, a heuristic tool introduced, during the first ten years of the twentieth century, by P. Weiss and L.S. Ornstein in the classical theory of phase transitions. Yet, the approximation is not acceptable here without some further discussion since it leads to the paradoxes already mentioned.

We are now in a position to recognize Haag's seminal contribution [Haag, 1962]: the *tacit* assumption of the irreducibility of the representation of the field algebra is *untenable*. Giving up this assumption allows one to resolve the paradoxes: Δ and hence the coefficients u and v in the Bogoliubov–Valatin transformation (77) — rather than being multiples of the identity — now belong to the *non-trivial* center \mathcal{Z} of the representation canonically associated by the GNS construction corresponding to the equilibrium state of the system. The gauge group now acts in a non-trivial manner on \mathcal{Z} and thus restores the symmetry of the theory. And in

the limit considered, the time-evolution is well defined as an automorphism group of the von Neumann algebra generated by the representation. These technical niceties have been successively refined — and confirmed — in subsequent investigations, cf. e.g. [Emch and Guenin, 1966; Thirring and Wehrl, 1967; Thirring, 1968; Dubin and Sewell, 1970; Sewell, 1982b].

5.2 Beyond Fock space: The Bose gas

Even before the modular structures were formally recognized by mathematicians, their first instantiation appeared in QSP. One can indeed discern these structures in the pioneering re-examination Araki and Woods [1963] made of the Bose–Einstein model for an ideal quantum gas; for the original version of the model, cf. subsection 2.4. The present subsection summarizes the principal aspects of the Araki–Woods treatment.

The reader is assumed to be familiar with the definition of the Weyl form of the canonical commutation relations (CCR) for a countably infinite number of degrees of freedom, as a family $\{W(f) \mid f \in \mathcal{D}(\mathbb{R}^3)\}$ of unitary operators acting on the (boson) Fock space $\mathcal{F} := \bigoplus_{N=0}^{\infty} {}^s\mathcal{H}^N$ and satisfying $\forall f, g \in \mathcal{D}(\mathbb{R}^3) : W(f)W(g) = \exp\{-i \operatorname{Im}(f, g)/2\}$; where $\mathcal{D}(\mathbb{R}^3)$ is the space of all infinitely differentiable functions $f : \mathbb{R}^3 \rightarrow \mathbb{C}$ which have compact support; and ${}^s\mathcal{H}^N$ is the symmetric N -fold tensor product of the one-particle space $\mathcal{H}^1 = \mathcal{L}^2(\mathbb{R}^3)$ with itself; cf. e.g. [Emch, 1972a], or [Halvorson, 2006].

For the Bose gas at temperatures $T > T_c$ where T_c is the critical temperature found by Bose and Einstein, the GNS representation π_g corresponding to the gaseous normal phase — in the thermodynamical limit at fixed density ρ and chemical activity z — is given as follows. The Hilbert space of the representation π_g may be identified with $\mathcal{H} = \mathcal{F} \otimes \mathcal{F}$; its cyclic vector is $\Phi = \Phi_o \otimes \Phi_o$, where Φ_o is the vacuum vector in \mathcal{F} . Then

$$(80) \quad \pi_g[W(f)] = W(\zeta_+f) \otimes W(K\zeta_-f)$$

where completeness demands that we specify that $(\zeta_+f)^\sim(k) = [1 + \rho(\beta, z; k)]^{\frac{1}{2}} \tilde{f}(k)$, and $(\zeta_-f)^\sim(k) = [\rho(\beta, z; k)]^{\frac{1}{2}} \tilde{f}(k)$, $(Kf)^\sim(k) = \tilde{f}(k)^*$; $\rho(\beta, z; k) = z[\exp(\beta\epsilon(k)) - z]^{-1}$ with $\epsilon(k) = |k|^2/2m$ and z is determined by $\rho = (2\pi)^{-3} \int d^3k \rho(\beta, z; k)$.

The von Neumann algebra $\mathcal{N}_g = \{\pi_g[W(f)] \mid f \in \mathcal{D}(\mathbb{R}^3)\}''$ is a factor, the commutant of which is $\mathcal{N}'_g = \{\nu_g[W(f)] \mid f \in \mathcal{D}(\mathbb{R}^3)\}'$ where

$$(81) \quad \nu_g[W(f)] = W(K\zeta_-f) \otimes W(\zeta_+f).$$

Note that ν_g also gives a representation of the Weyl CCR.

In what I believe was the first presentation of the programme proposed in [Haag *et al.*, 1967] to a wide audience of mainstream physicists, namely the huge IU-PAP 1966 Copenhagen meeting on statistical mechanics, Winnink [Winnink, 1967] started indeed with a summary of the above results. As the duality between the

von Neumann algebra and its commutant is already a property of finite systems — see scholium 33 and remark 34 above — Winnink’s emphasis was that this property may persist in general for systems endowed with infinitely many degrees of freedom, as is the case in this specific model — the Bose gas — where the thermodynamical limit of canonical equilibrium is controlled. The emphasis on dealing with infinite systems — also advocated in the lecture [Verboven, 1967] preceding Winnink’s — raised eyebrows with many of the physicists in the Copenhagen audience, to wit: “Wouldn’t one think that, so to say, the motivation of going to an infinite system would be to obtain simpler results than are obtained for a finite system?” [Uhlenbeck, 1967]; or even more pointedly: “What does this have to do with statistical mechanics?” [van Kampen, 1967]. The conjecture was already floated that the formalism could be useful for an adequate description of phase transitions, a conjecture I will examine in subsections 5.6 and 5.7.

In retrospect, it is quite remarkable that Araki and Woods [1963] had already unearthed several features that were later placed in the context of the general theory that was to be built on the subsequent work of Tomita and Takesaki [Tomita, 1967; Takesaki, 1970a] for the mathematical formalism and the work of Haag *et al.* [1967] for its application to QSP. Among the results by Araki and Woods, one may note that the von Neumann factor \mathcal{N}_g they constructed for $T > T_c$ is of type III — a type of factor the existence of which was known, but for which examples were then quite elusive even in the pure mathematics literature — and this was the first occurrence of this type of factor in QSP, although their ubiquity was later recognized all over in QSP and in QFT; and also in pure mathematics, but that is another story. In addition, Araki and Woods established that the unitary operators implementing time-evolution and space-translation on the von Neumann algebra \mathcal{N}_g do not belong to this algebra. They also discussed the representations relative to the superfluid phase which occurs for temperatures $0 < T < T_c$, and they found that the associated GNS representation is an integral of factor representations. Incidentally, they do mention that this points to a formal analogy with the mathematical structure Haag found in his study of the BCS model; see subsection 5.1 above.

5.3 The KMS condition and the Heisenberg model

The first proof that the KMS conditions themselves are actually satisfied in concrete infinite quantum systems was provided by Araki [1969] for a class of one-dimensional quantum spin-lattice models which includes the archetypal model — originally proposed by Heisenberg [1928] as a putative model for ferromagnetism — defined by the local, so-called “exchange” Hamiltonian:

$$(82) \quad H_\Lambda = -J \sum_{k=a}^{b-1} \sigma_k \cdot \sigma_{k+1}$$

where J is the coupling constant describing interactions of neighbouring quantum spins $\sigma_k = (\sigma_k^x, \sigma_k^y, \sigma_k^z)$ sitting on a regular, one-dimensional finite string $\Lambda =$

$[a, b] \subset Z^1$; and $\sigma_k \cdot \sigma_{k+1}$ denotes $\sigma_k^x \sigma_{k+1}^x + \sigma_k^y \sigma_{k+1}^y + \sigma_k^z \sigma_{k+1}^z$.

The problem of determining whether this quantum model would support ferromagnetism in its thermodynamical limit — even in this one-dimensional version — turned out to be much harder to handle than the classical Ising model where only interactions $J\sigma_k^z \sigma_{k+1}^z$ are considered.

For the classical models, a method known as the *transfer-matrix* — and actually proposed for the two-dimensional Ising model [Kramers and Wannier, 1941] — allows one to treat the one-dimensional version of this classical model in a few lines for nearest neighbour interactions, or even with strictly finite-range interactions, i.e. when the interactions are strictly zero between spins that are further apart than a finite distance (the same for all pairs). At the cost of quite some work [Ruelle, 1968b], the method can be made to work for interactions that extend to infinity, while decaying sufficiently fast so as to have finite moment or so that the surface energy has a bound independent of the volume.

As emphasized in some detail in [Emch, 1972b], even the nearest neighbour quantum Heisenberg model requires an extension akin to the method used for the infinite range classical case. Araki [1969] thus managed to control the thermodynamical limit φ of the canonical equilibrium state and its time correlation functions well enough to establish that for all positive temperatures $\beta > 0$, the state φ satisfies the KMS condition; and that it is extremal with respect to this condition — i.e. cannot be decomposed into a mixture of states satisfying the KMS condition — and shows no spontaneous magnetization. Thus, while physicists — with their understanding of the onset of collective behaviour in the classical case — could anticipate that this quantum one-dimensional model would not exhibit any ferromagnetic phase transition, Araki proves it.

The class of models for which Araki established the above results is strongly dependent on the one-dimensionality of the “lattice” \mathbb{Z} . Nevertheless, the proof does not require that the interactions be isotropic, i.e. couplings between the different components of the spins do not need to be the same in all directions. Moreover the proof does not require that the interactions between the spins on the lattice be restricted to nearest neighbours: in the original version of the proof, it was only required that they vanish between spins that are further apart than a fixed (but arbitrary) finite distance, but even this restriction has been relaxed to cover the same range as the corresponding classical models. Finally, whereas in the Heisenberg model the individual half-spins are described by copies of the algebra $\mathcal{M}(2, \mathbb{C})$, the proof accommodates as well the cases where $\mathcal{A}_k \simeq \mathcal{M}(n, \mathbb{C})$ with $n < \infty$.

Hence Araki’s results support a conjecture that pure thermodynamical phases may be described by extremal KMS states; see subsection 5.7 below for further evidences.

It may be added that if, in addition, lattice translation invariance is built into the theory through the local Hamiltonian H_Λ , e.g. as in (82), then the limiting KMS state φ is invariant under the group \mathbb{Z} of the lattice translations, and — since its GNS representation leads to a factor — φ is also extremal with respect to this

condition, so that space-correlations between spins decay very fast as their distance increases. Here, technically speaking, φ is exponentially, uniformly clustering in space; i.e. for any quasi-local observable $A \in \mathcal{A}_o$, there exist positive constants γ and δ such that for all finite N and all $B \in \mathcal{A}_{\mathbb{Z} \setminus [-N, N]}$: $|\varphi(AB) - \varphi(A)\varphi(B)| \leq \delta \|B\| \exp(-\gamma N)$.

Before leaving the Heisenberg model, note that in the case $T = 0$, it also offers a very nice toy model for QFT; cf. e.g. [Streater, 1967].

5.4 *The KMS condition and stability*

The following five points summarize the KMS story I have told so far.

1. Von Neumann's definition of quantum canonical states at finite temperature is limited to finite systems; and this limitation renders cumbersome (at best) the formalism's application to QSP — see subsection 3.3.
2. For finite systems, the von Neumann equilibrium states are exactly those that satisfy a formal analytic condition, the so-called KMS condition — see subsection 4.1.
3. The KMS condition can be extended beyond the mathematical formalism laid down by von Neumann — see subsection 4.2.
4. The KMS condition is satisfied in some concrete models of infinite systems for states that have a reliable interpretation as temperature equilibrium states — see subsections 5.2 and 5.3.
5. The KMS condition appears as well in a purely mathematical context, the Tomita–Takesaki theory of modular algebras which turns out to be very fertile; while the latter aspect of the story would carry us beyond the bounds of this essay, some hints are briefly alluded to in see subsections 4.2 and 5.2.

Before the KMS theory could be deemed adequate as a physical theory, it ought to meet at least two more concerns: (i) the formalism should allow the mathematical description of quantum phenomena that escape the grip of von Neumann's formalism; (ii) the KMS states should be stable. Subsections 5.2 and 5.3 above indicate how the first of these two concerns is met; further examples will be presented in subsection 5.7. The present subsection addresses the second concern, as various stability criteria — labeled A to E — are discussed. The order of the presentation is to direct the reader's attention to the progressive emergence of formulations whereby KMS states are characterized in terms increasingly germane to those of variational principles.

A. Cut-and-paste stability.

We begin with a model that is sufficiently simple to provide exact results supporting the expectation that a large system in a canonical equilibrium state should be

able to serve as a *thermal reservoir* for “any” of its parts. The model is a variation on the theme of the so-called XY-model; this variation was proposed and solved in [Emch and Radin, 1971]; further references will be given at the end of this subsection.

The X-Y model itself — referred to below as the ‘un-partitioned system’ — is a one-dimensional quantum spin-lattice gas with finite-range interactions. Specifically, for any region $\Lambda = \{k \in \mathbb{Z} \mid a \leq k \leq b\}$ with $-\infty < a + 1 < b < \infty$, the Hamiltonian is

$$(83) \quad H_{[a,b]} = - \sum_{k=a}^{b-1} (1 + \zeta) \sigma_k^x \sigma_{k+1}^x + (1 - \zeta) \sigma_k^y \sigma_{k+1}^y \quad .$$

From the work of Araki — see subsection 5.3 above — we learned that the thermodynamical limit (as $a \rightarrow -\infty, b \rightarrow \infty$) of both of the following objects exists: (i) the canonical equilibrium state $\varphi_{[a,b]}$ at any finite natural temperature $\beta > 0$ and (ii) the time-evolution $\alpha_{[a,b]}$; and that the resulting state φ and evolution α of the infinite system satisfy the KMS condition.

We now split the total system in two non-interacting parts: a finite region to which we affix the subscript S , and an infinite region to which we affix the subscript R , which is the complement of Λ_S in \mathbb{Z} , namely:

$$\Lambda_S = [c, d] \quad \text{and} \quad \Lambda_R = (-\infty, c - 1] \cup [d + 1, \infty)$$

with $-\infty < a < c - 1; c < d - 1; d + 1 < b < \infty$.

This partitioned system can be viewed as the thermodynamical limit of a finite system with Hamiltonian:

$$(84) \quad \tilde{H}_{[a,b]} = H_{[a,c-1]} + H_{[c,d]} + H_{[d+1,b]} \quad .$$

Clearly the C^* -algebras for both the original system and the partitioned system are the same, namely the C^* -inductive limit $\mathcal{A} := \otimes_{k \in \mathbb{Z}} \mathcal{A}_k$ where the \mathcal{A}_k are copies of the algebra $\mathcal{M}(2, \mathbb{C})$ of the 2×2 matrices with complex entries. Thus $\mathcal{A} = \mathcal{A}_S \otimes \mathcal{A}_R$ where $\mathcal{A}_S := \otimes_{k \in \Lambda_S} \mathcal{A}_k$ and $\mathcal{A}_R := \otimes_{k \in \Lambda_R} \mathcal{A}_k$.

Again, as for the original (un-partitioned) system, the thermodynamical limit of the canonical equilibrium state and of the evolution of the partitioned system, defined from (84), do exist and satisfy the KMS condition; they are denoted here by $\tilde{\varphi}$ and $\tilde{\alpha}$.

Note that φ and $\tilde{\varphi}$ are different. For instance, φ is invariant with respect to translations along the chain, while $\tilde{\varphi}$ is not. Nevertheless, a first stability property of this model is established in [Emch and Radin, 1971], namely:

$$(85) \quad \forall A \in \mathcal{A} : \lim_{|t| \rightarrow \infty} \tilde{\varphi}(\alpha_t[A]) = \varphi(A) \quad .$$

Hence, as the full evolution α unfolds, the correlations between S and R that were cut by the partitioning are re-established: the partition is erased.

Moreover, let $\tilde{\varphi}_S$ denote the restriction of $\tilde{\varphi}$ to \mathcal{A}_S ; and similarly for R ; one has $\tilde{\varphi} = \tilde{\varphi}_S \otimes \tilde{\varphi}_R$. The evolution $\tilde{\alpha}$ preserves the partitioning, namely $\forall A \in$

\mathcal{A}_S [resp. \mathcal{A}_R] : $\tilde{\alpha}_t[A] \in \mathcal{A}_S$ [resp. \mathcal{A}_R]. Hence, we have $\tilde{\alpha} = \tilde{\alpha}_S \otimes \tilde{\alpha}_R$; i.e. the two systems evolve independently. Again the KMS conditions are satisfied for S and R separately.

After this partitioning, let us now change the temperatures of S and R so that (i) $\tilde{\varphi}_{S,\beta_S}$ is the canonical equilibrium on \mathcal{A}_S at some temperature β_S with respect to the evolution $\tilde{\alpha}_S$; and similarly (ii) with β_R (possibly different from β_S) for $\tilde{\varphi}_{R,\beta_R}$ (w.r.t $\tilde{\alpha}_R$) on \mathcal{A}_R . Let thus $\tilde{\varphi}_{S,\beta_S} \otimes \tilde{\varphi}_{R,\beta_R}$ be the initial state of the partitioned system; and denote by φ_β the canonical equilibrium on the whole system, at temperature β with respect to the original interacting evolution α . Then, the following is proven [Emch and Radin, 1971] for all $\beta_S, \beta_R > 0$ and for all $A \in \mathcal{A}$:

$$(86) \quad \lim_{|t| \rightarrow \infty} \tilde{\varphi}_{S,\beta_S} \otimes \tilde{\varphi}_{R,\beta_R}(\alpha_t[A]) = \varphi_\beta(A) \quad \text{with } \beta = \beta_R .$$

Hence the name ‘*cut-and-paste stability*’. The system is first partitioned in two parts, a finite system S surrounded by an infinite system R that do not interact with one another: the interactions across the boundaries [i.e. between the sites $c-1$ and c ; and between the sites d and $d+1$] have been ‘cut’. In this configuration, the finite system S and the infinite system R are separately put at (different) temperatures β_S and β_R . When these systems are ‘pasted’ back together, one finds that the joint evolution drives the full system $S \cup R$ to a temperature β which has to be the temperature at which R was initially, namely $\beta = \beta_R$. In this sense, R serves as a thermal reservoir for S .

The special property of the model responsible for this result is that it satisfies a remarkable condition which I now describe.

Let γ be the automorphism of \mathcal{A} uniquely determined by

$$(87) \quad \forall k \in Z : \begin{cases} \gamma[\sigma_k^z] = & \sigma_k^z \\ \gamma[\sigma_k^x] = - & \sigma_k^x \\ \gamma[\sigma_k^y] = - & \sigma_k^y \end{cases} .$$

Note in particular that the original Hamiltonian $H_{[a,b]}$ and the cut Hamiltonian $\tilde{H}_{[a,b]}$ belong to the even algebra $\mathcal{A}_e := \{A \in \mathcal{A} \mid \gamma[A] = A\}$. This entails that in the thermodynamical limit $\varphi \circ \gamma = \varphi$ and $\gamma \circ \alpha \circ \gamma = \alpha$; and similarly for all the corresponding objects obtained after the partitioning. In particular, the evolution preserves the even algebra, i.e. $\forall (t, A) \in \mathbb{R} \times \mathcal{A}_e : \alpha_t[A] \in \mathcal{A}_e$.

Now, the special property of the model can be explicitly stated:

$$(88) \quad \forall A, B \in \mathcal{A}_e : \lim_{|t| \rightarrow \infty} \|A\alpha_t[B] - \alpha_t[B]A\| = 0 \quad ;$$

i.e. the evolution, when restricted to the even observables of the model, is *strongly asymptotically abelian*.

The proof — and an immediate generalization — of (86) above is a direct consequence of quantum ergodic theory (see subsection 3.5). First, one notices that φ is uniformly clustering in space, i.e. for every $\epsilon > 0$ and every $A \in \mathcal{A}$ there exists

a finite region Λ such that for every B outside this region $|\varphi(AB) - \varphi(A)\varphi(B)| \leq \epsilon\|B\|$. This entails that the KMS state φ is extremal with respect to this condition, i.e. cannot be decomposed into a convex combination of other KMS states; see subsection 5.6 below, and in particular definition 57. These properties are inherited by the restriction φ_e of φ to the even algebra \mathcal{A}_e . The asymptotic abelianness of the evolution then implies [Araki and Miyata, 1968] that the state φ_e is not only time-invariant — as we know every KMS state must be — but it is also extremal with respect to this condition, i.e. cannot be decomposed into other time-invariant states, which is to say that φ_e cannot be written as $\varphi_e = \lambda\psi_e + (1 - \lambda)\chi_e$ with $0 < \lambda < 1$ and ψ_e, χ_e time-invariant, *unless* $\psi_e = \chi_e = \varphi_e$.

Since $\varphi, \varphi_e, \tilde{\varphi}_S, \tilde{\varphi}_R$ are even, one loses no information by carrying out the proof of (85) and (86) in \mathcal{A}_e ; in particular, (88) implies the existence of the point-wise limit of $(\varphi_S \otimes \varphi_R) \circ \alpha_t$ in the LHS of (86); then the above argument shows that it must coincide with φ .

Note further that what is proven in [Emch and Radin, 1971] is in fact a stronger result, which implies (86) and thus (85) as particular cases, namely that for all even states ψ_S of the system S :

$$(89) \quad \forall A \in \mathcal{A} : \lim_{|t| \rightarrow \infty} \psi_S \otimes \tilde{\varphi}_R(\alpha_t[A]) = \varphi_\beta(A)$$

which therefore reinforces the ‘cut-and-paste stability’ interpretation proposed immediately after equation (86).

This result may be further generalized in two ways. First, the restriction that ψ_S in (89) be an even state can be dispensed with; cf. [Araki and Barouch, 1982]. Second, as was already noticed in [Emch and Radin, 1971], an ergodic or averaged version of (89), specifically, with η denoting an invariant mean on the group \mathbb{R} :

$$(90) \quad \forall A \in \mathcal{A} : \eta\{\psi_S \otimes \tilde{\varphi}_R(\alpha[A])\} = \varphi_\beta(A)$$

obtains [Emch and Radin, 1971], even if only a weaker version of (88) holds, namely the condition of η -asymptotic abelianness (60), i.e.

$$(91) \quad \forall A, B, C \in \mathcal{A}_e : \eta\{\varphi(C^* [A\alpha[B] - \alpha[B]A]C)\} = 0 \quad .$$

Depending on one’s intellectual temperament, *either* the general argument presented earlier, *or* the specific model just reviewed, raises the question of whether the KMS condition could be derived from some general stability argument. This question is addressed from several angles in (B)–(E) below.

The model presented above was discussed again in [Robinson, 1973]; see also [Araki and Barouch, 1982] and references therein. It belongs to a long line of inquiries that started as attempts to derive Newton’s cooling law from first principles; related problems are usually understood under the generic label ‘return to equilibrium’. As of this writing, the latest comprehensive entry on the subject may be [Bach *et al.*, 2000] where a wealth of “novel technical devices” are brought to bear; the reader will also find there an informative sample of the large literature on

the subject. In a broad sense, several — but not all — of the criteria of stability in this subsection also address this perennial problem of return to equilibrium from small or local deviations. Its ubiquity however should not overshadow two other important and largely unsolved problems; cf. subsection .6.4 below.

B. Stability against local perturbations.

Various conditions of asymptotic abelianness were investigated by Kastler *et al.* For a summary, see [Kastler, 1976] which also offers a annotated bibliography. For their main stability theorem, they settled on the notions described in definitons 41 and 42 below.

DEFINITION 41. Let \mathcal{A} be a C^* -algebra. An evolution $\alpha : \mathbb{R} \rightarrow \text{Aut}(\mathcal{A})$ is said to be L^1 - asymptotically abelian on a norm dense $*$ -subalgebra $\mathcal{A}_o \subset \mathcal{A}$ when

$$\forall (t, A) \in \mathbb{R} \times \mathcal{A}_o : \alpha_t[A] \in \mathcal{A}_o ;$$

and

$$\forall A, B \in \mathcal{A}_o : \int_{-\infty}^{+\infty} dt \|B \alpha_t[A] - \alpha_t[A] B\| < \infty .$$

Some preliminary notations are required for Definition 42 below. With \mathcal{A} and α as in definition 41 let $\mathcal{A}_{sa} = \{A \in \mathcal{A} \mid A = A^*\}$, and let \mathcal{S} be the set of all states on \mathcal{A} , equipped with its weak topology. For $\varphi \in \mathcal{S}$ and an element $h \in \mathcal{A}_{sa}$ with $\varphi(h^2) > 0$, define

- (i) the perturbed state φ^h by $\varphi^h : A \in \mathcal{A} \rightarrow \frac{1}{\varphi(h^2)} \varphi(hAh) \in \mathcal{A}$;
- (ii) the perturbed evolution $\{\alpha^h_t \mid t \in \mathbb{R}\}$ by $\alpha^h_t : A \in \mathcal{A} \mapsto U^h_t \alpha_t[A] U^{h*}_t$ where $\{U^h_t \mid t \in \mathbb{R}\}$ satisfies the so-called ‘co-cycle differential equation’ (the derivative is w.r.t. the norm-topology)

$$\forall t \in \mathbb{R} : i \frac{d}{dt} U^h_t = U^h_t \alpha_t[h] \quad \text{with initial condition } U^h_o = I .$$

To understand the sense in which α^h may be viewed as the perturbed evolution corresponding to h , note that the above co-cycle equation admits a unique continuous solution $t \in \mathbb{R} \mapsto U^h_t \in \mathcal{A}$; it can be computed explicitly as the norm-convergent Dyson series:

$$U^h_t = \sum_{n=0}^{\infty} C^h_{t,n} \quad \text{with } C^h_{t,n} = (-i)^n \int_o^t dt_n \int_o^{t_n} dt_{n-1} \dots \int_o^{t_2} dt_1 \alpha_{t_1}[h] \dots \alpha_{t_n}[h].$$

This solution satisfies: (i) U^h_t unitary, and

$$(ii) \quad \forall s, t \in \mathbb{R} : U^h_{s+t} = U^h_s \alpha_s[U^h_t] .$$

Consequently, the evolution defined as $\{\alpha^h_t \mid t \in \mathbb{R}\}$ is a group of automorphisms of \mathcal{A} with, in particular:

$$\forall s, t \in \mathbb{R} : \alpha^h_{s+t} = \alpha^h_s \circ \alpha^h_t .$$

The interpretation of α^h as the evolution resulting from the perturbation of α by the operator h obtains from the following relation between the generators of α^h and α :

$$i \frac{d}{dt} \alpha^h_t \Big|_{t=0} = i \frac{d}{dt} \alpha_t \Big|_{t=0} + \delta^h \quad \text{with} \quad \delta^h : A \in \mathcal{A} \mapsto [h, A] := hA - Ah \in \mathcal{A} \quad .$$

DEFINITION 42. With the above notations, an α -invariant state φ on \mathcal{A} is said to be stable against inner perturbations, whenever there is a neighbourhood $\mathcal{V}_\varphi \subset \mathcal{S}$ of φ such that $\forall A \in \mathcal{A}$, and $\forall h \in \mathcal{A}_{sa}$ with $\varphi^h \in \mathcal{V}_\varphi$:

1. $\forall t \in \mathbb{R} : \varphi^h(\alpha^h_t[A]) = \varphi^h(A)$;
2. with $\lambda \in \mathbb{R} : \lim_{\lambda \rightarrow 0} \varphi^{\lambda h}(A) = \varphi(A)$;
3. $\lim_{t \rightarrow \infty} \varphi^h(\alpha_t[A]) = \varphi(A)$.

THEOREM 43. With \mathcal{A} , \mathcal{A}_o and α as in definition 41, assume that α is L^1 -asymptotically abelian on \mathcal{A}_o . Let φ be an α -invariant state on \mathcal{A} and assume that φ is stable against inner perturbations in the sense of definition 42. Then — under three ancillary conditions to be discussed below — φ satisfies the KMS condition with respect to α for some natural temperature β .

REMARKS 44. The ancillary conditions of the theorem are sketched in the three entries below.

1. The state φ is assumed *not* to be a trace, i.e. there exist $A, B \in \mathcal{A}$ such that $\varphi(AB) \neq \varphi(BA)$. This is meant to avoid the classical circumstance that would arise in the limit of infinite temperature, i.e. $\beta = 0$, i.e. $T = \infty$.
2. In the GNS representation canonically associated to φ , the generator of the unitary group $U(\mathbb{R})$ that implements $\alpha(\mathbb{R})$ is assumed *not* to be one-sided. This is meant to avoid the opposite circumstance where φ would be a zero-temperature ground state, i.e. $\beta = \infty$, i.e. $T = 0$.
3. The state φ is assumed to be hyperclustering of order 4 on the $*$ -subalgebra \mathcal{A}_o . This technical condition requires the following to hold: for every positive integer $p \leq 4$ and all $A_1, \dots, A_p \in \mathcal{A}_o$, there exist positive constants C and δ such that

$$(92) \quad \forall t_1, \dots, t_p \in \mathbb{R} : \varphi_p^T(\alpha_{t_1}[A_1] \cdots \alpha_{t_p}[A_p]) \leq C \{ 1 + \max |t_k - t_l|^{1+\delta} \}^{-1}$$

where the truncated correlations φ_p^T are defined recursively by $0 = \varphi_o^T$, $\varphi(A) = \varphi_1^T(A)$ and $\varphi(A_1, \dots, A_p) = \sum_P \varphi_{n_1}^T(A_{k_1}, \dots, A_{k_{n_1}}) \cdots \varphi_{n_j}^T(A_{q_1}, \dots, A_{q_{n_j}})$ and the sum carries over all order-preserving partitions of $S = \{1, 2, \dots, p\}$ in subsets $S_j \subseteq S$ satisfying the following conditions: $S = \cup_j S_j$, $j \neq k \Rightarrow S_j \cap S_k = \emptyset$, and within each $S_j = \{k_1, k_2, \dots, k_{n_j}\} : k_1 < k_2 < \dots < k_{n_j}$. The reader will verify immediately that $\varphi(A_1, A_2) =$

$\varphi_2^T(A_1, A_2) + \varphi_1^T(A_1)\varphi_1^T(A_2)$, and then realize that the recursion relation explains better what is going on with higher truncated correlations than writing explicitly the summations over P .

Note that the φ_p^T provide a hierarchy where all correlations of lower order already have been taken into account. In particular in the case of the CCR, a remarkable result of Robinson [Robinson, 1965] shows that either this hierarchy goes up indefinitely or, if the truncated φ_n^T vanish for all $n \geq N$ with $N > 2$, then they must vanish for all $n > 2$.

The concept of truncated φ_n^T is not a stranger. It comes to us as a quantum cousin of the “cumulants” of classical probability theory and of the “Ursell functions” of classical statistical mechanics. The classical equivalent of Robinson’s theorem gives a characterization of the Gaussian distribution, which translates in quantum statistics as yet another characterization of the canonical equilibrium state of an assembly of free harmonic oscillators. Robinson’s theorem thus gives a foretaste of why it is so difficult to produce and/or control models of QFT and QSP that are not “quasi-free”.

To sum up, the third ancillary condition of the theorem aims to convey that in the course of time all time-correlations of order $p \leq 4$ are to decay rapidly enough for long time separations.

The investigations by Kastler *et al.* reported above appear to be systematically predicated on conditions of time-asymptotic abelianness (definition 41) and time-hyperclustering (remark 44(3)). Thus compare these with any of the conditions encountered in sections 3 and 4; the latter are naturally satisfied for space translations, but in constructing specific models, even these conditions are extremely difficult to impose straight on the microscopic dynamics, i.e. on the Hamiltonian that is to describe the time evolution. Whether this is an intrinsic shortcoming of the theory behind theorem 43 above, or an indication of some lack of either imagination or technical dexterity on the part of model builders remains open at this stage. Nevertheless, it appears that one weak form of asymptotic abelianness is not only sufficient but also necessary when one wants to identify, among KMS states, those that are merely extremal with respect to this condition, from those that are, moreover, extremal with respect to time-invariance; cf. e.g. [Emch, 1972a, corollary 2, p. 206]; or remark 63(6) below. Here again, reminiscences from the perennial ergodic dreams in classical statistical *Hamiltonian mechanics* would incline some to hope that such an identification could perhaps be in the cards. As I have recognized in several other parts of this essay, my crystal ball remains clouded on this issue.

C. Thermal reservoir stability.

Consider the intuitive idea that a system R may be construed as a “thermal reservoir” at temperature β , if it drives suitably devised test systems S to equilibrium at temperature β when they are coupled to R . Kossakowski *et al.* [1977] proposed to formalize this idea in the following manner; see also [Sewell, 2002, pp. 114–116].

For a concrete motivation, compare the specific XY-model described in part **A** of this subsection.

To model situations where one expects that R ought to be very much larger than S in order to exclude feedbacks from the test system S onto the reservoir R , one assumes that R is infinite and S is finite.

The putative reservoir R is described by a triple $\{\mathcal{A}^R, \alpha^R, \varphi^R\}$ where \mathcal{A}^R is a C^* -algebra; α^R is an evolution group of automorphisms of \mathcal{A}^R ; and φ^R is a state on \mathcal{A}^R , invariant under the evolution α^R . Denote by $\delta^R := i \frac{d}{dt} \alpha^R_t \Big|_{t=0}$ the generator of the evolution α^R . Some ancillary conditions on R will be specified later.

The test system S is a dynamical system in the sense of von Neumann, i.e. is described by: $\{\mathcal{A}^S, \alpha^S, \varphi^S\}$ where $\mathcal{A}^S = \mathcal{B}(\mathcal{H})$; α^S is the evolution generated by a Hamiltonian H^S such that for all temperatures $\beta > 0$, $Z := \text{Tr} \exp(-\beta H^S) < \infty$; and φ^S is given by

$$(93) \quad \varphi^S_\beta(A^S) = \text{Tr} \rho^S_\beta A^S \quad \text{with} \quad \rho^S_\beta = Z^{-1} e^{-\beta H^S} .$$

$\delta^S = i [H^S, \cdot]$ will denote the generator of α^S . Finally, \mathfrak{S}^S will denote the set of all countably additive states on \mathcal{A}^S .

A family $\{\alpha^\lambda \mid \lambda \geq 0\}$ of dynamical couplings between R and S is described by groups of automorphisms on $\mathcal{A} = \mathcal{A}^R \otimes \mathcal{A}^S$, the generator of which is of the form:

$$(94) \quad \left. \begin{aligned} \delta^\lambda &= \delta^R \otimes I + I \otimes \delta^S + \lambda \delta_V \quad \text{where} \\ \delta_V : A \in \mathcal{A} &\mapsto i \lambda [V, A] \in \mathcal{A}, \quad \text{with} \quad V \in \mathcal{A}_{sa} \end{aligned} \right\} .$$

As the ancillary conditions on R are specified, so will be the form of V ; see (97) and (98) below.

The next step in the modeling is devised to emphasize the sense in which the long-time cumulative effects on S of the evolution α^λ are accounted for when R and S are coupled. For this Kossakowski *et al.* [1977] appeal to the so-called van Hove limit, an instance of which already appeared in subsection 3.5; see also remark 45 below. For the system at hand here, the van Hove limit takes the following form. First, it considers only a reduced evolution, namely only what the system S experiences of the total evolution; mathematically this reduction is achieved by $E : \mathcal{A} \rightarrow \mathcal{A}_S$, the conditional expectation defined, for all $A_S \otimes A_R$, by $E[A_S \otimes A_R] = A_S \varphi_R(A_R)$, and then extended by linearity and continuity to \mathcal{A} . Secondly, the van Hove limit requires to focus on a long-time/weak-coupling regime defined by rescaling time with an inverse power of the interaction strength. Thus, the van Hove limiting procedure consists here in proving that the following limit exists for all positive ‘rescaled’ times s :

$$(95) \quad \gamma^S_s : A_S \in \mathcal{A}^S \mapsto \lim_{\substack{\lambda \rightarrow 0; t \rightarrow \infty \\ s = \lambda^2 t}} \alpha^S_{-t} \circ E \circ \alpha^\lambda_t[A_S] \in \mathcal{A}^S .$$

REMARKS 45. This type of limit has a long history. I learned it first from van Hove [van Hove, 1955] where the author had proposed it as a tool to relate

macroscopic transport phenomena to the microscopic dynamics that is expected to underlie them. It emphasizes that in such discussions time ought to be rescaled in a way determined by the strength λ of the interaction. Some justifications for taking such a limit will be discussed in subsection 6.1 below.

Finally, given two C^* -algebras \mathcal{A} and \mathcal{B} , and n a non-negative integer, one says that a map $\gamma : \mathcal{A} \rightarrow \mathcal{B}$ is *n-positive* whenever it is linear, and the induced map $\gamma_n : \mathcal{A} \otimes \mathcal{M}(n, \mathbb{C}) \rightarrow \mathcal{B} \otimes \mathcal{M}(n, \mathbb{C})$ is positive, i.e. the image of any positive element in $\mathcal{A} \otimes \mathcal{M}(n, \mathbb{C})$ is a positive element in $\mathcal{B} \otimes \mathcal{M}(n, \mathbb{C})$. When either \mathcal{A} or \mathcal{B} is abelian, a positive map is necessarily n -positive; hence n -positivity is a notion new to the non-commutative context of QSP. Furthermore, a map is said to be *completely positive* whenever it is n -positive for all $n \in \mathbb{Z}^+$. In connection with expressions like the right-hand side of (95) above, note that the composition of completely positive maps is again completely positive; and that automorphisms, states, injections and conditional expectations are completely positive maps. A collection $\{\gamma_s \mid s \in \mathbb{R}^+\}$ of maps of \mathcal{A} into itself is said to form a semi-group whenever γ_0 is the identity map, and $\forall (s, t) \in \mathbb{R}^+ \times \mathbb{R}^+ : \gamma_{s+t} = \gamma_s \circ \gamma_t$.

This should exhaust the list of general preliminaries necessary to describe the stability criterion proposed by Kossakowski *et al.* [1977], namely:

DEFINITION 46. A system $\{\mathcal{A}^R, \alpha^R, \varphi^R\}$ is said to be a thermal reservoir at temperature β whenever there is a “large enough” collection \mathfrak{T}_β of test systems $\{\mathcal{A}^S, \alpha^S, \varphi_\beta^S\}$ and dynamical couplings $\{\alpha^\lambda\}$ such that

1. the van Hove limit (95) exists, and defines a semi-group of completely positive transformations $\{\gamma_s^S \mid s \in \mathbb{R}^+\}$ of \mathcal{A}_S ;
 2. the canonical von Neumann equilibrium state φ_β^S on \mathcal{A}_S is the only state $\varphi \in \mathfrak{S}^S$ that is invariant under both α^S and γ^S ;
- (96) 3. $\forall (\psi^S, A^S) \in \mathfrak{S}^S \otimes \mathcal{A}^S : \lim_{s \rightarrow \infty} \psi^S(\gamma_s^S[A_S]) = \varphi_\beta^S(A_S) \quad .$

The term “large enough” in the above definition admittedly needs to be made more precise: this is where the ancillary conditions on the interaction V and the reservoir R enter the picture and allow one to prove scholium 47 and theorem 48 below.

One condition is that the interaction V in (94) be of the form

$$(97) \quad V = \sum_{k=1}^n B_k^R \otimes B_k^S \quad \text{with} \quad \left\{ \begin{array}{l} n \text{ is finite} \\ B_k^R \in \mathcal{A}_{sa}^R \text{ and } \varphi^R(B_k^R) = 0 \\ B_k^S \in \mathcal{A}_{sa}^S \end{array} \right\} \quad .$$

Note that the conditional expectation $E[V]$ of V vanishes.

An additional condition is that there exists $\mathcal{A}_o^R \subseteq \mathcal{A}^R$ such that: (i) $\text{Span}\{\mathcal{A}_o^R \cup I\}$ (where I is the identity in \mathcal{A}^R) is norm dense in \mathcal{A}^R ; (ii) for all $B_k^R \in \mathcal{A}_o^R$, the functions $t \mapsto \varphi^R(B_j^R \alpha_t^R[B_k^R])$ are in L^1 ; and (iii) the multi-time truncated correlations, for the state φ^R to be tested, satisfy

$$(98) \quad t_1 < \dots < t_l \text{ with } |t_j - t_k| \rightarrow \infty \Rightarrow \{\varphi^R\}^T (\alpha_{t_1}^R[B_{t_1}^R] \dots \alpha_{t_n}^R[B_{t_n}^R]) \rightarrow 0.$$

Upon taking advantage of [Davies, 1974, theorem 2.3], the following results were obtained in [Kossakowski *et al.*, 1977]:

SCHOLIUM 47. These ancillary conditions are sufficient to imply that condition (1) in definition 46 is satisfied for all finite S .

This ensures that the collection \mathfrak{T}_β of test systems will indeed be large enough.

THEOREM 48. *When the circumstances just outlined are realized, the following conditions are equivalent:*

1. *for some temperature β , the state φ^R is a KMS state on \mathcal{A}_R with respect to the evolution α_R ;*
2. *the system R , in the state φ^R , is a thermal reservoir for temperature β in the sense of definition 46 with “large enough” sharpened by scholium 47.*

REMARKS 49.

1. Hence, every test system S in \mathfrak{T}_β is driven to equilibrium at temperature β by the reservoir R exactly when φ^R satisfies the KMS condition for this temperature.
2. As mentioned before, this result is largely model-independent, and does not involve, *at least explicitly*, any assumption of time asymptotic abelianness. Moreover, instead of a single model for which one can prove that a special infinite system in equilibrium serves as thermal reservoir for each of its finite parts, the present theorem characterizes a collection \mathfrak{T}_β of test systems S for which the infinite system R serves as a thermal reservoir. Thus, the theorem is an improvement on the particular motivating model discussed in paragraph **A** above.
3. Yet, as [Kossakowski *et al.*, 1977] noticed, the decay of multi-time correlations (98) that enables their proposal to work is reminiscent of the similar conditions imposed by Kastler *et al.* in their result on stability against local perturbations; see (92) in paragraph **B** above.
4. From an empirical point of view, the theorem may be regarded as specifying a procedure to lift the notion of temperature in equilibrium QSP from finite systems to infinite systems.
5. Nevertheless, it must be noted that the circumstances under which γ^S is known to satisfy condition (1) of definition 46 and the decay of correlations in (98) do appear to involve some clustering properties that may limit the domain of applicability of the theorem to cases where $\pi_\varphi(\mathcal{A}_R)''$ is a factor, and where φ^R is extremal under both the KMS condition and the condition

of time-invariance. Hence, asymptotic abelianness enters less conspicuously here.

D. Passivity.

In [Pusz and Woronowicz, 1978] the authors noticed a property of KMS states which they called *passivity*; and they found ways to show that this property in turn entails the KMS property under assumptions that *do not* involve asymptotic abelianness in *time*.

Specifically, let $\{\mathcal{A}, \varphi, \alpha\}$ be a dynamical system where \mathcal{A} is a C^* -algebra, φ be a state on \mathcal{A} , and $\{\alpha_t \mid t \in \mathbb{R}\}$ is a one-parameter group of automorphisms of \mathcal{A} . Let then $D(\delta)$ denote the domain of the generator δ of the evolution α , i.e. $D(\delta)$ is the linear subspace of all $A \in \mathcal{A}$ such that the derivative $\delta[A] := i \frac{d}{dt} \alpha_t[A]$ exists.

Consider now the situation obtained by letting this system interact during a finite time-interval with an outside system, so that the effect of their interaction on the system of interest may be assumed to be described as the perturbed dynamics α^h satisfying the differential equations:

$$(99) \quad \forall A \in D(\delta) : \begin{cases} i \frac{d}{dt} \alpha^h_t[A] = \alpha^h_t[\delta[A] + [h_t, A]] \\ \alpha^h_{t=0}[A] = A \end{cases}$$

where h is an element of $C^1_+(\mathbb{R}, \mathcal{A}_{sa})$, the set of all continuously differentiable functions, with compact support in \mathbb{R}^+ and taking their values in the self-adjoint part of \mathcal{A} . The system is thus an open system for all times t in the support of h , i.e. for all times when the perturbation h is actually in effect. The condition that the support of h be compact and contained in \mathbb{R}^+ ensures that, for all times $T > \sup \{t \in \mathbb{R} \mid h_t \neq 0\}$, the external conditons are as they were at time $t = 0$. The smoothness condition $h \in C^1$ on the time-dependence of the external pertubation is a mathematical convenience that is physically reasonable. Then

$$(100) \quad L^h_T(\varphi) := \int_0^T dt \varphi(\alpha^h_t [\frac{d}{dt} h_t])$$

describes the energy transmitted to the system in the time interval $[0, T]$ during which the system was under the influence of the external perturbation h .

DEFINITION 50. The state φ is said to be passive if for all $h \in C^1_+(\mathbb{R}, \mathcal{A}_{sa})$ and all $T > \sup \{t \in \mathbb{R} \mid h_t \neq 0\} : L^h_T(\varphi) \geq 0$.

Upon having advanced this definition, Pusz and Woronowicz [1978] proved the following result:

THEOREM 51. Let $\{\mathcal{A}, \varphi, \alpha\}$ be a C^* - dynamical system, and consider the following conditions: (I) φ is either a KMS state with respect to α for some temperature $\beta > 0$; or is a ground state; and (II) φ is passive in the sense of definition 50. Then:

1. Without further assumptions: (I) \Rightarrow (II) .

2. If furthermore: (i) \mathcal{A} admits an action $\nu : G \rightarrow \text{Aut}(\mathcal{A})$ where G is a locally compact amenable group; (ii) ν commutes with the evolution α , i.e. $\forall (t, g) \in \mathbb{R} \times G : \nu_g \circ \alpha_t = \alpha_t \circ \nu_g$; and (iii) φ is η -clustering with respect to the action of G . Then these conditions, taken together, entail $(II) \Rightarrow (I)$.

REMARKS 52. The following remarks focus on part (2) of the theorem, i.e. the operational characterization of KMS states as passive.

1. In the passivity condition (II), φ has not been assumed to be invariant under the unperturbed evolution α ; in part (2) this property obtains as φ is proven to satisfy the KMS condition.
2. The condition that ν commutes with α is natural in view of the conclusion to be obtained: if an automorphism leaves invariant a KMS state, then it must commute with the evolution with respect to which this state is KMS.
3. Invariant means and amenable groups were introduced in subsection 3.5; see in particular definition 21 and remark 31(5).
4. Among the ancillary assumptions listed in (2), it is not even necessary to impose as a precondition that φ be G -invariant; this follows from the explicit assumption that it is η -clustering, i.e. (see definition 22):

$$\forall A, B \in \mathcal{A} : \eta^G(\varphi(\nu_g[A] B)) = \varphi(A) \varphi(B) \quad .$$

Actually, this condition entails furthermore that φ cannot be decomposed in a convex combination of other G -invariant states.

5. In QSP, the natural candidate for G is the group of translations in *space*. Hence, in contrast with the stability conditions studied earlier, the assumed clustering property does not need to be with respect to time. This allows us to consider systems for which the evolution is *not* asymptotically abelian. This opening is significant when it comes to concrete modeling for the purposes of QSP: one may not wish to have to identify the weak-clustering with respect to the group $\{\nu_g \mid g \in G\}$ and any putative clustering with respect to the evolution $\{\alpha_t \mid t \in \mathbb{R}\}$.
6. In addition, Pusz and Woronowicz [Pusz and Woronowicz, 1978] propose an alternative route, replacing all the ancillary conditions in part (2) of the theorem by a strengthened form of passivity. Specifically, instead of considering a single dynamical system, they consider, for every positive integer, identical non-interacting copies $\{\mathcal{A}_k, \varphi_k, \alpha_k \mid k = 1, \dots, N\}$ from which one constructs the collective dynamical system $\{\mathcal{A}^N, \varphi^N, \alpha^N\}$ where $\{\mathcal{A}^N = \otimes_{k=1}^N \mathcal{A}_k, \varphi^N = \otimes_{k=1}^N \varphi_k, \text{ and } \alpha^N = \otimes_{k=1}^N \alpha_k\}$. The perturbation h however is allowed to be a general element in $C_+^1(\mathbb{R}, \mathcal{A}^N)$, so that α^h is allowed not to act independently on each of the component systems. Then φ is said to be *completely passive* whenever for every positive integer N the

state φ^N is passive. Now, without further ado — i.e. without having to impose condition (2) in theorem 51 — the complete passivity of φ can be proven to be equivalent to the condition that φ satisfy the KMS condition. For QSP, the choice between the condition of complete passivity or condition (2) in the theorem, is largely a question of taste.

E. Thermodynamical stability.

To close this subsection, I wish to indicate how the concept of thermodynamical stability gives rise to yet another characterization of KMS states, this one without restriction on whether the states considered are to be extremal with respect to the KMS condition. To avoid technicalities, I present these considerations in the simplest case, namely where the system is a quantum spin-lattice and thus is described by a C^* -algebra $\mathcal{A} = \otimes_{k \in \mathbb{Z}^d} \mathcal{A}_k$ where the \mathcal{A}_k are copies of a finite matrix algebra, say $\mathcal{M}(n, C)$, with n and d finite. Throughout $\Lambda \subset \mathbb{Z}^d$ denotes a connected finite subset of the lattice \mathbb{Z}^d ; φ denotes a state on \mathcal{A} ; φ_Λ denotes the restriction of φ to the finite matrix algebra $\mathcal{A}_\Lambda = \otimes_{k \in \Lambda} \mathcal{A}_k$; and ρ_Λ is the density matrix corresponding to φ_Λ . Furthermore it is convenient to assume here that the dynamics obtains from short-range — or possibly suitably tempered — interactions between the sites. The reader interested in how far the considerations presented below may be pursued will find a review in [Sewell, 2002]; among the original papers, let it suffice to mention for orientation purposes [Araki, 1974; Araki and Sewell, 1977; Sewell, 1977; Sewell, 1980b; Ruelle, 1968a; Robinson, 1971; Araki and Moriya, 2002].

A version of the second law of thermodynamics — compare with the equivalent form of the variational principle defined immediately after theorem 3 — defines the local free-energy relative to Λ at natural temperature $\beta = 1/kT$ as:

$$F_{\Lambda, \beta}(\varphi) = E_\Lambda(\varphi) - T S_\Lambda(\varphi) \quad \text{with} \quad \begin{cases} E_\Lambda(\varphi) &= \varphi_\Lambda(H_\Lambda) \\ S_\Lambda(\varphi) &= -k \text{Tr } \rho_\Lambda \log \rho_\Lambda \end{cases} .$$

Two states ψ and φ on \mathcal{A} are said to satisfy the equivalence relation $\overset{\Lambda_o}{\sim}$ whenever they coincide outside the finite region Λ_o . We then write $\psi \sim \varphi$ whenever there exists Λ_o such that $\psi \overset{\Lambda_o}{\sim} \varphi$. For the quantum lattice considered here, one can then prove that the following limit exists

$$(101) \quad \forall \psi \sim \varphi : \Delta F_\beta(\psi | \varphi) := \lim_{\Lambda \uparrow \mathbb{Z}^d} (F_{\Lambda, \beta}(\psi) - F_{\Lambda, \beta}(\varphi)) .$$

For the order of the arguments ψ and φ in ΔF_β recall that mathematicians (and some philosophers) read from right to left, while most physicists seem to read from left to right. Thus, $\Delta F_\beta(\psi | \varphi)$, as written above, represents the increment of free-energy when passing from the state φ to any state ψ that differs from φ only in a finite region. Araki and Sewell [Araki and Sewell, 1977; Sewell, 1977] introduced the following definition and prove the following result; see also [Sewell, 1980b; Sewell, 2002].

DEFINITION 53. With $\Delta F_\beta(\varphi | \psi)$ as in (101), a state φ on \mathcal{A} is said to be locally thermodynamically stable at natural temperature β whenever

$$\forall \psi \sim \varphi : \Delta F_\beta(\psi | \varphi) \geq 0 \quad .$$

Hence, to require that this stability condition be satisfied is indeed a variational principle: the free-energy of the state φ cannot be reduced by going to a state ψ that differs from φ only locally.

THEOREM 54. *For a state φ on a quantum lattice system of the type considered here, the following conditions are equivalent:*

1. φ satisfies the KMS condition at natural temperature β ;
2. φ is locally thermodynamically stable at natural temperature β .

REMARKS 55.

1. This result involves *in an essential manner* the local structure of the system considered, namely that the global algebra \mathcal{A} is the C^* -inductive limit of local algebras \mathcal{A}_Λ relative to bounded regions of space, where the indexing net $\mathcal{F} := \{\Lambda\}$ is absorbing, i.e. — recall subsection 3.4, part I — for every point x in space, there is a bounded region $\Lambda \in \mathcal{F}$ such that $x \in \Lambda$. An alternative version is requiring that for every bounded region Ω of space there exists some $\Lambda \in \mathcal{F}$ such that $\Omega \subseteq \Lambda$; both versions are acceptable in axiomatic QSP.
2. Also in contrast to reservoir stability — see theorem 48 — this result is one of *internal* consistency in the sense that it establishes the equivalence of two definitions of equilibrium for the *same* system when described from two different points of view: the microscopic KMS condition and the local aspect of thermodynamics of the system considered. In particular, the argument does not involve any coupling of the system considered with any test system.
3. Extensions of the domain of validity of the theorem are desirable. In this respect, quantum spin-lattice systems with reasonably long-range interactions have been controlled. However, some technical difficulties often stand in the way toward the expected extensions to continuous systems. Typically these difficulties originate in the infinite dimensionality of the Hilbert spaces \mathcal{H}_Λ corresponding to finite regions, and in the fact that the corresponding Hamiltonians H_Λ are unbounded; also, precautions may have to be taken to ensure that the local particle-density remains bounded.
4. One type of extension of the above variational principle is instructive, namely the shift from local stability to global stability requirements. Specifically, consider again a quantum-lattice system defined on \mathbb{Z}^d . Assume further that the dynamics is invariant under the translation group $G = \mathbb{Z}^d$, and restrict attention to the set \mathfrak{S}^G of states ψ each of which is G -invariant. Assume finally that the following limits exist

$$(102) \quad f_\beta(\psi) = \lim_{\Lambda \uparrow \mathbb{Z}^d} |\Lambda|^{-1} F_{\Lambda, \beta}(\psi) \quad ; \quad \phi_\beta = \lim_{\Lambda \uparrow \mathbb{Z}^d} |\Lambda|^{-1} \log \text{Tr} \exp^{-\beta H_\Lambda} \quad .$$

A state $\varphi \in \mathfrak{S}^G$ is now said to be *globally thermodynamically stable* — or GTS for short — whenever it minimizes the free-energy density, i.e. when

$$(103) \quad f_\beta(\varphi) = \min_{\psi \in \mathfrak{S}^G} f_\beta(\psi) = \phi_\beta \quad .$$

As long as one remains with G –invariant states on quantum-lattices having G –invariant dynamics with only *short-range* interactions, one has

$$\varphi \text{ is GTS} \iff \varphi \text{ is KMS} \quad .$$

However, while \Rightarrow remains valid even when interactions are allowed to extend over a reasonably long range, the “short-range” requirement is essential for \Leftarrow . It has been suggested [Sewell, 1980b] that KMS states that are not GTS, i.e. do not minimize the free-energy density, may model metastable states.

5.5 A brief excursion into QFT

As a remark on the role of KMS states in mathematical physics I wish to mention, however briefly, the appearance of modular structures beyond the confines of non-relativistic QSP, namely their entry into relativistic QFT. For the general framework of algebraic QFT, cf. e.g. in this volume [Halvorson, 2006]; for a presentation specifically geared to QFT on curved space-times, cf. also [Wald, 1994]; and for a discussion of some of the interpretation problems raised by the materials in this section, cf. [Clifton and Halvorson, 2001].

From the perspective developed in this essay, the natural entry into the considerations to be discussed in the present subsection is through a manifestation, in Minkowski-space QFT, of the Tomita–Takesaki duality — recall scholium 33 or theorem 39.

Bisognano and Wichmann [Bisognano and Wichmann, 1975] developed a consequence of a standard result in axiomatic QFT — the Reeh–Schlieder theorem, cf. e.g. [Streater and Wightman, 1964, p.168], or [Emch, 1972a, p. 290] and references cited therein — which ensures in particular that the vacuum state φ , when restricted to a wedge $W_R = \{(x, y, z, t) \in M^{3+1} \mid z > |t|\}$, is faithful on the corresponding algebra \mathcal{N}_R . Thus, this restriction φ_R of φ to \mathcal{N}_R equips the latter with the structure of a Tomita–Takesaki modular algebra. Here, the canonical objects of the Tomita–Takesaki theory have a seminal geometric interpretation. The involutive antiunitary operator J — corresponding to the reflection $(x, y, z, t) \rightarrow (x, y, -z, -t)$ which maps the wedge W_R to the wedge $W_L = \{(x, y, z, t) \in M^{3+1} \mid z < |t|\}$ — implements a bijection from \mathcal{N}_R to $\mathcal{N}'_R \simeq \mathcal{N}_L$; and the modular group $\{\Delta^{i\lambda} \mid \lambda \in \mathbb{R}\}$ implements on \mathcal{N}_R the Lorentz

boost

$$\begin{pmatrix} x \\ y \\ z \\ t \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cosh(2\pi\lambda) & -\sinh(2\pi\lambda) \\ 0 & 0 & -\sinh(2\pi\lambda) & \cosh(2\pi\lambda) \end{pmatrix} \begin{pmatrix} x \\ y \\ z \\ t \end{pmatrix} .$$

Since uniformly accelerated observers moving in the interior of a wedge W_R perceive its boundaries as past and future horizons, the result of Bisognano and Wichman could be interpreted as saying that in the universe of such observers — the wedge W_R — the vacuum of the field in M^{3+1} appears to be a thermal bath, in the following sense. The state $\varphi_R : N \in \mathcal{N}_R \mapsto \varphi(N) \in \mathbb{C}$ — where $\mathcal{N}_R \subset \mathcal{N}$ is the algebra corresponding to the wedge W_R , \mathcal{N} is the algebra corresponding to the full Minkowski space, and φ is the vacuum defined on \mathcal{N} — is a KMS state at temperature $\beta > 0$ with respect to the evolution $\{\tau_t : N \in \mathcal{N}_R \mapsto \tau_t[N] = \Delta^{-it/\beta} N \Delta^{it/\beta} \in \mathcal{N}_R \mid t \in \mathbb{R}\}$ (where, as usual, the numerical value of the natural temperature $\beta = 1/kT$ depends on the scale with respect to which the time t is measured).

The physical interest of this interpretation is enhanced by an earlier remark by Rindler [Rindler, 1966] to the effect that the universe of uniformly accelerated observers in W_R is similar to the universe around the Schwarzschild solution of the Einstein equations, i.e. around a stationary “black hole”.

With this dictionary in hand, the phenomenon discovered by Bisognano and Wichmann as a consequence of the Tomita–Takesaki theory translates into an effect found independently by Unruh [1976] in an attempt to clarify the then recently discovered Hawking effect [Hawking, 1975] (also known as the Hawking radiation). The latter describes a related but different phenomenon, the creation of thermally distributed particles around a collapsing black hole. The similarities and differences between the Unruh and the Hawking effects are discussed in [Wald, 1994, chapters 5 and 7]; for some of the thermodynamical aspects of the subject in the astrophysics literature, cf. e.g. [Davies, 1978; Hawking and Page, 1983] or [Wald, 1994, chapter 6]; for the specific questions of *what* is actually measurable, *how* and *where*, see [Unruh and Wald, 1984]; and for some of the philosophical issues, see [Clifton and Halvorson, 2001].

The impact of the Bisognano and Wichmann discovery on the axiomatic QFT literature began with the work of Sewell [Sewell, 1980a; Sewell, 1982a] who generalized their results to some curved manifolds, perceived the role that the bifurcate horizon plays in the Unruh effect, and proposed to identify the Hawking temperature and the temperature in the KMS condition associated with the Tomita–Takesaki modular theory. The introduction of KMS structures in QFT then turned out to be the harbinger of a “*revolution*” [Borchers, 2000]. A few among the many developments that ensued are: an extension of the axiomatic of algebraic QFT to curved manifolds; an interpretation of the intrinsic geometry of space-time in terms of consistency relations between the modular structures to be attached to an absorbing net of intersecting wedge-like regions of GR

space-times; and the beginnings of a relativistic QSP where local KMS conditions are formulated in terms of future-directed time-like vectors that determine local rest-frames; cf. e.g. [Summers and Verch, 1996; Buchholz *et al.*, 2002; Ojima, 2003; Wiesbrock, 1997; Buchholz, 2003; Summers and White, 2003; Buchholz and Lechner, 2004]; closer to the Hawking effect proper, [Haag *et al.*, 1994; Kay and Wald, 1991; Fredenhagen and Haag, 1990]; and for a new framework [Fredenhagen, 2003].

5.6 A mathematical interlude: extremal KMS states

The role of extremal KMS states in QSP will be studied in subsection 5.7. The purpose of this section is to review some mathematical preliminaries such as the definition of extremal KMS states, their characterization in terms of their GNS representation, and the decomposition of a KMS state into its extremal components.

REMARKS 56.

1. Let \mathcal{A} be a C^* -algebra, $\beta > 0$ and τ be a group of automorphisms of \mathcal{A} . The set \mathfrak{S}_β of all KMS states on \mathcal{A} that satisfy the KMS condition for τ and β is convex, i.e. for any two KMS states ψ and χ on \mathcal{A} , with respect to the same τ and β , and any $\lambda \in (0, 1)$: $\varphi = \lambda\psi + (1 - \lambda)\chi$ is again a KMS state for τ and β .
2. The set \mathfrak{S}_β is closed in the w^* - topology it inherits from \mathcal{A} , and it is bounded in the metric topology. Hence it is w^* -compact, and the Krein-Milman theorem entails that \mathfrak{S}_β is the w^* -closed convex hull of the set \mathfrak{E}_β of its extreme points [Dunford and Schwartz, 1964, theorem V.8.4]. This ensures not only the existence of extremal points, but also that there are sufficiently many of them: every element in \mathfrak{S}_β is the limit of finite convex sums of elements in \mathfrak{E}_β ; see definition 57 below.
3. Moreover $\beta_1 \neq \beta_2$ entails $\mathfrak{S}_{\beta_1} \cap \mathfrak{S}_{\beta_2} = \emptyset$. Incidentally, the GNS representations constructed from states $\varphi_1 \in \mathfrak{S}_{\beta_1}$ and $\varphi_2 \in \mathfrak{S}_{\beta_2}$ with $\beta_1 \neq \beta_2$ are disjoint in the sense that no subrepresentation of one of these is unitarily equivalent to any subrepresentation of the other; cf. [Takesaki, 1970c].

DEFINITION 57. Given a von Neumann algebra \mathcal{N} , a group $\{\tau_t \mid t \in \mathbb{R}\}$ of automorphisms of \mathcal{N} , $\beta \in \mathbb{R}^+$, and \mathfrak{S}_β as in remark 56(1) above. A state $\varphi \in \mathfrak{S}_\beta$ is said to be extremal KMS at natural temperature β if it does not admit a convex decomposition into states in \mathfrak{S}_β — i.e. states that satisfy the KMS condition for the same τ and β . The set of all extremal KMS states is denoted \mathfrak{E}_β .

THEOREM 58. Let φ be a faithful normal state on a von Neumann algebra \mathcal{N} and τ be the unique group of automorphisms of \mathcal{N} with respect to which φ satisfies the KMS condition for some natural temperature β . Denote by \mathcal{Z} the center $\mathcal{N} \cap \mathcal{N}'$ of \mathcal{N} . Then

A. For every $(t, Z) \in \mathbb{R} \times \mathcal{Z}$, $\tau_t[Z] = Z$.

B. For every positive non-zero element $Z \in \mathcal{Z}$ with $0 < Z < I$,

$$\psi(N) := \varphi(Z)^{-1}\varphi(ZN) \quad \text{and} \quad \chi(N) := \varphi(I - Z)^{-1}\varphi((I - Z)N)$$

define two states ψ and χ on \mathcal{N} that satisfy the KMS condition for the same τ and β and provide a convex decomposition of φ .

C. For every φ that admits a convex decomposition $\varphi = \lambda\psi + (1 - \lambda)\chi$ into states ψ and χ on \mathcal{N} that satisfy the KMS condition for the same τ and β , there exists a unique positive non-zero element $Z \in \mathcal{Z}$ with $\|Z\| \leq 1$ such that for all $N \in \mathcal{N}$

$$\psi(N) = \varphi(Z)^{-1}\varphi(ZN) \quad \text{and} \quad \chi(N) = \varphi(I - Z)^{-1}\varphi((I - Z)N).$$

Proof. As pointed out in remark 37, we may assume without loss of generality that \mathcal{N} is presented in standard form, so that there exists a cyclic and separating vector $\Phi \in \mathcal{H}$ for \mathcal{N} with $\forall N \in \mathcal{N} : (\Phi, N\Phi) = \varphi(N)$.

[A.] $Z \in \mathcal{Z} \Rightarrow \forall (t, N) \in \mathbb{R} \times \mathcal{N}$, $\varphi(N^*\tau_t[Z]) = \varphi(\tau_t[z]N^*)$ and thus φ being KMS entails that $\varphi(N\tau_t[Z])$ is constant in t so that $\forall t \in \mathbb{R} : (N\Phi, [\tau_t[Z] - Z]\Phi) = 0$. Φ being cyclic entails $[\tau_t[Z] - Z]\Phi = 0$, and then Φ being separating entails $[\tau_t[Z] - Z] = 0$.

[B.] φ being faithful and $0 < Z < I$ positive and non-zero entail $0 < \varphi(Z) < 1$; and, upon taking into account that Z and thus $Z^{\frac{1}{2}}$ belong to \mathcal{N}' , one verifies that ψ and χ are states on \mathcal{N} and that they inherit from φ its KMS property. Moreover, one reads immediately from their definition that $\varphi = \lambda\psi(N) + (1 - \lambda)\chi(N)$, where $0 < \lambda = \varphi(Z) < 1$.

[C.] Conversely, from $\varphi = \lambda\psi(N) + (1 - \lambda)\chi(N)$, with $0 < \lambda < 1$ one has $\psi \leq \lambda^{-1}\varphi$ and thus there exists an element $X \in \mathcal{N}'$ such $\forall N \in \mathcal{N} : \psi(N) = (X\Phi, NX\Phi)$, i.e. ψ is a vector state on \mathcal{N} and thus is normal and majorized by the normal functional $\lambda^{-1}\varphi$. Hence the Sakai-Radon-Nikodym [Sakai, 1971, proposition 1.24.4], entails that there exists some positive $Y \in \mathcal{N}$ with $\|Y\| \leq 1$ such that

$$\forall N \in \mathcal{N} : \psi(N) = \frac{1}{2}\lambda^{-1}\varphi(NY + YN).$$

Suppose that there exists another element $\tilde{Y} \in \mathcal{N}$ with the same properties. Let then $X = Y - \tilde{Y}$. We have then $0 = \varphi(X^*X + XX^*)$ and thus, since φ is a positive linear functional and both X^*X and XX^* are positive: $\varphi(X^*X) = 0$. Since φ is faithful, $X = 0$ i.e. $Y = \tilde{Y}$ i.e. Y is unique.

It remains to be shown that the assumptions of the theorem entail that Y also belongs to \mathcal{N}' . Since φ and ψ are KMS, they satisfy for all $t \in \mathbb{R} : \varphi \circ \tau_t = \varphi$ and $\psi \circ \tau_t = \psi$. Consequently

$$\psi(N) = \psi(\tau_t[N]) = \frac{1}{2}\lambda^{-1}\varphi(\tau_t[N]Y + Y\tau_t[N]) = \frac{1}{2}\lambda^{-1}\varphi(N\tau_{-t}[Y] + \tau_{-t}[Y]N).$$

From the uniqueness of $Y \in \mathcal{N}$ which we just established, we have $\forall t \in \mathbb{R} : \tau_t[Y] = Y$. φ being KMS entails therefore $\forall N \in \mathcal{N} : \varphi(NY) = \varphi(YN)$ and thus $\psi(N) = \lambda^{-1}\varphi(YN)$. Upon applying the KMS condition to both ψ and φ , we get $\forall N \in \mathcal{N} : NY = YN$ i.e. $Y \in \mathcal{N}'$. Clearly then $\lambda = \varphi(Z)$. The same argument goes through with χ replacing ψ and $(I - Z)$ replacing Z . ■

The following characterization is an immediate consequence of the above theorem:

COROLLARY 59. *With the assumptions of theorem 58, the KMS state φ is extremal KMS iff \mathcal{N} is a factor, i.e. iff \mathcal{N} has trivial center: $\mathcal{N} \cap \mathcal{N}' = \mathbb{C}I$.*

SCHOLIUM 60. With the assumptions of theorem 58, assume that φ is not extremal KMS, but that the center \mathcal{Z} of \mathcal{N} is generated by a family $\{P_k \in \mathcal{Z} \mid k = 1, 2, \dots\}$ of mutually orthogonal projectors. Then there exists a unique decomposition of φ into a convex combination $\sum_k \lambda_k \varphi_k$ of states φ_k on \mathcal{N} where the φ_k are extremal KMS for the same dynamics τ and the same natural temperature β .

Proof. To say that φ is a KMS state that is not extremal KMS is to say that there exist KMS states ψ_j and scalars $\mu_j \in (0, 1)$ such that $\varphi = \sum_j \mu_j \psi_j$. From part C of the theorem, for every ψ_j there exists a positive $Z_j \in \mathcal{Z}$ such that $\forall N \in \mathcal{N} : \psi_j = \phi(Z_j)^{-1}\phi(Z_j N)$. Since \mathcal{Z} is an abelian von Neumann algebra with discrete spectrum, every Z_j may be written as $\sum_k z_k P_k$ with $z_k \in \mathbb{R}^+$ and the P_k are minimal projectors in \mathcal{Z} . Hence the $\varphi_k : \mathcal{N} \in \mathcal{N} \mapsto \lambda_k^{-1} \varphi(P_k N) \in \mathbb{C}$ with $\lambda_k = \phi(P_k)$ are states on \mathcal{N} . From part B of the theorem, these are still KMS states for the same τ and β . Therefore, it only remains to show that the states φ_k are extremal with respect to the KMS condition.

To see this, consider the decomposition $\mathcal{H} = \oplus_k \mathcal{H}_k$ where \mathcal{H}_k are the subspaces $\{\Psi_H \in \mathcal{H} \mid P_k \Psi = \Psi\}$. Since each P_k belongs to \mathcal{Z} , the subspaces \mathcal{H}_k are stable under \mathcal{N} and under \mathcal{N}' , i.e. whenever $X \in \mathcal{N}$ or $X \in \mathcal{N}'$, we have $\forall \Psi \in \mathcal{H}_k : X\Psi \in \mathcal{H}_k$. Let then $\mathcal{N}_k = \{P_k N P_k \mid N \in \mathcal{N}\}$, $\mathcal{N}'_k = \{P_k N P_k \mid N \in \mathcal{N}'\}$; and note that these are von Neumann algebras acting on the space \mathcal{H}_k admitting there a cyclic and separating vector, namely $P_k \Phi$, such that $\forall N \in \mathcal{N}_k : \tilde{\varphi}_k(N) := (\Phi_k, N \Phi_k)$ defines a faithful normal state on \mathcal{N}_k ; it is thus the restriction to this algebra of the state φ . Note further that for all $t \in \mathbb{R}$, \mathcal{N}_k is stable under τ_t . Since $\mathcal{N}_k \cap \mathcal{N}'_k = \mathbb{C}I_k$ (where I_k is the identity operator in \mathcal{H}_k) $\tilde{\varphi}_k$ is extremal KMS. Proceeding *ab absurdo*, suppose that φ_k itself is not extremal KMS. Then there would exist some KMS state ψ on \mathcal{N} and some $\lambda \in (0, 1)$ such that $\psi \leq \lambda^{-1}\varphi_k$. Denote by $\tilde{\psi}_k$ the restriction of ψ to \mathcal{N}_k . We have then, in particular, $\lambda^{-1}\varphi_k(N^* N) \geq \psi([N P_k]^* [N P_k]) = \psi_k(N^* N)$; i.e. $\lambda^{-1}\tilde{\varphi}_k \geq \tilde{\psi}_k$. Since $\tilde{\varphi}_k$ is extremal KMS and $\tilde{\psi}_k$ is KMS, the equality must prevail, i.e. $\lambda^{-1}\tilde{\varphi}_k = \tilde{\psi}_k$; and since $\tilde{\varphi}_k$ and $\tilde{\psi}_k$ are states, $\lambda = 1$, i.e. on $(\mathcal{N}_k)_{\tilde{\varphi}_k} : \tilde{\psi}_k(N^* N) = \tilde{\varphi}_k(N^* N)$. By the Schwartz inequality, we have for every $N \in \mathcal{N}$, $\tilde{\psi}_k([P_k N P_k]^* [P_k N P_k]) \leq \psi(N^* N)$, and thus $\psi \geq \varphi_k$. Together with the initial inequality, namely $\psi \leq \varphi_k$ (since we know now that $\lambda = 1$), these two inequalities reduce to $\psi = \varphi_k$. Hence φ_k is indeed an extremal KMS state on \mathcal{N} . Since φ_k is extremal KMS on \mathcal{N}_k , the restriction

$\tilde{\psi}_k$ of ψ to this algebra must coincide with $\tilde{\varphi}_k$; and thus φ_k is maximal KMS on \mathcal{N} . Hence φ has been decomposed into a convex combination of extremal KMS states. Uniqueness follows by contradiction. ■

DEFINITION 61. A convex set \mathcal{C} is said to be a simplex whenever every point in \mathcal{C} admits a unique convex decomposition into extremal points of \mathcal{C} .

Recall that in two-dimensional Euclidean geometry, a triangle is a simplex; indeed any point in the triangle obtains as a unique convex combination of points situated at the vertices of the triangle. But a circle is not a simplex: the set of its extreme points is the circumference of the circle, and given any point inside the circle, all secants through this point give different convex combinations of extreme points.

REMARKS 62.

1. Scholium 60 may therefore be paraphrased by saying that under the assumption that the spectrum $Sp(\mathcal{Z})$ of the center \mathcal{Z} is discrete, \mathfrak{S}_β is a simplex; and that the decomposition is a weighed sum with respect to a discrete probability measure supported by the extreme points \mathfrak{E}_β of the set \mathfrak{S}_β of all normal KMS states for the given dynamics τ and the given natural temperature β . From the proof of the scholium, one checks that the latter statement extends indeed to all normal states, and not just to those that are faithful.
2. In case $Sp(\mathcal{Z})$ is not discrete, the above sum must be replaced by an integral, and some measure-theoretical trimmings are necessary to specify the sense in which φ defines a *unique* measure concentrated on the boundary of \mathfrak{S}_β . The general mathematical context in which these decompositions appear is in the study of central measures, see [Takesaki, 1970a; Kadison and Ringrose, 1983/1986]. For the purpose of this essay, the simpler version just described will suffice to anchor the conceptual structure of the theory governing the unique decomposition of KMS states into their extremal components.
3. Note that the set of states on a quantum system described by the von Neumann postulates is *not* a simplex: if a density matrix has at least one eigenvalue with multiplicity greater or equal to 2, its decomposition in pure states is not unique. Hence the set of KMS states in *quantum* statistical physics possesses a *classical* property that is otherwise not heard of in the quantum realm.
4. It remains to demonstrate that this property is relevant to QSP; and thus that QSP requires the consideration of situations for which the relevant representations do not lead to factors, in contrast with the von Neumann formalism of quantum mechanics where the canonical equilibrium states lead only to factor representations — recall the end of remark 34. This problem is the object of the next subsection.

5.7 Extremal KMS states, pure thermodynamical phases

The main argument one can advance to justify the claim that pure thermodynamical phases be described in QSP as extremal KMS states originates in the conjunction of three circumstances.

The *first* is based on subsections 5.3 and 5.4 where strong evidences were given for the identification of canonical equilibrium states as KMS states.

The *second* is the fact that extremal KMS states are the elementary objects in the KMS theory. This is reminiscent of the identification of atomic levels in *spectroscopy* with irreducible representations of the group of symmetries of the system, the famous “Gruppenpest” epitomized in [Wigner, 1931]. In mathematics, this programme was extended to a systematic presentation of the familiar so-called *special functions* where these functions now appear as bases of irreducible representations of groups; cf. e.g. [Talman, 1968; Vilenkin, 1968]. Closer to the focus of this essay, the early identification — in [Murray and von Neumann, 1936] — of *factors as the building blocks* of the theory of von Neumann algebras proceeds from the same principle: a methodological option confirmed by the central decomposition of a von Neumann algebra as a direct integral of factors; cf. e.g. [Kadison and Ringrose, 1983/1986, theorem 14.2.2, pp. 1027–1028]. All the while, the group-theoretical approach continues to contribute in sorting out qualitative classification problems in *nuclear spectroscopy*, and *elementary particles* high energy physics.

The *third* circumstance pointing to the description of pure thermodynamical phases as extremal KMS states — i.e. KMS states the GNS representations of which are factors, cf. corollary 59 above — is the mathematical fact that the decomposition of a KMS state in extremal KMS states is *unique*; cf. scholium 60 and remark 62(2) above. In the context of QSP, this fact naturally directs attention to the situation encountered in thermodynamics where an equilibrium state decomposes uniquely into its pure thermodynamical phases.

Thus, this subsection is divided into two parts. In part **A**, the above speculations are confronted with a model for QSP where everything can be computed explicitly. In part **B**, the characterization of pure thermodynamical phases as extremal KMS states is brought to bear on a famous argument by Landau offering a fundamental microscopic distinction between solids and fluids in term of space correlation functions. This exemplifies how the unique decomposition of KMS states into their extremal components helps describe the coexistence of pure thermodynamical phases in QSP and spontaneous symmetry breaking in systems undergoing phase transitions. For further discussion of the latter, cf. [Liu and Emch, 2005].

A. Quantum Weiss–Ising models for ferromagnetism

Recall first the results of Araki reported in subsection 5.3 concerning the absence of a ferromagnetic phase transition: the unique KMS state for each of the models covered there is extremal.

To check how this coincidence fares with systems that *do* exhibit several ther-

modynamical phases, we turn now to a class of models that have a long history in the physics of phase transitions [Weiss, 1907; Brout, 1965], and are accepted by mathematicians to be amenable to a sufficiently rigorous treatment [Kac, 1968], namely the Weiss–Ising models for ferromagnetism.

Consider a one-dimensional lattice \mathbb{Z} where to every site $k \in \mathbb{Z}$ a quantum spin σ_k is attached. To every finite string $\Lambda \subset \mathbb{Z}$ is associated a Hamiltonian

$$(104) \quad H = - \sum_{k \in \Lambda} [B + B_{\Lambda,k}] \sigma_k^z \quad \text{with} \quad B_{\Lambda,k} = \frac{1}{2} \sum_{j \in \Lambda} J_{\Lambda,jk} \sigma_j^z$$

where B is interpreted as a homogeneous external magnetic field parallel to a fixed direction z ; and $B_{\Lambda,k}$ is an average magnetic field, the so-called “molecular” field, experienced by the spin at site k resulting from all other spins in the region Λ . The artificial assumption imposed on the models of the van der Waals or Weiss-type, which makes them exactly solvable in the thermodynamical limit, is that the strength of the interaction $J_{\Lambda,jk}$ decreases with the size $|\Lambda|$ of Λ ; compare this with the property $|v(p, q)| \leq c/|\Lambda|$ of the interaction of the BCS model in subsection 5.1.

Adopting here a simplified version of [Emch and Knops, 1970], we will assume that

$$(105) \quad J_{\Lambda,jk} = \begin{cases} |\Lambda|^{-1} J > 0 & \text{when } j \neq k \\ 0 & \text{otherwise} \end{cases} .$$

Upon controlling the thermodynamical limit $|\Lambda| \rightarrow \infty$, one finds that two extremal KMS states emerge when $T < T_c$ where $1/kT_c = \beta_c = J^{-1}$. These are recognized by the following properties of a global observable — cf. definition 15 and scholium 23 — namely, the magnetization \mathbf{M} , the three components of which

$$M^i = \text{weak op. limit}_{|\Lambda| \rightarrow \infty} \frac{1}{|\Lambda|} \sum_{k \in \Lambda} \sigma_k^i \quad (i = x, y, z)$$

are defined in the corresponding temperature-dependent representation. They satisfy

$$(106) \quad \begin{aligned} \text{(i)} \quad & M^x = M^y = 0 \\ \text{(ii)} \quad & M^z = \tanh[\beta(B + JM^z)] . \end{aligned}$$

REMARKS 63.

1. For the transverse components, M^x and M^y , (106.i) was to be expected from the symmetry of the system. The interesting part is the result for the component M^z parallel to the applied magnetic field: (106.ii) is the classical self-consistency equation: the model exhibits a phase transition as there exists a temperature T_c (with $\beta_c = J^{-1}$) below which M^z does not vanish as $B \rightarrow 0$, but tends to a finite, temperature-dependent value, the so-called spontaneous magnetization.

2. Hence in the thermodynamical limit, the problem of determining the extremal KMS states of the model admits two new solutions, absent above T_c ; these extremal KMS states exhibit the two opposite spontaneous magnetizations characteristic of the two pure thermodynamical phases familiar from the treatment of the classical case in e.g. [Kac, 1968].
3. A phase transition has occurred at $T = T_c$ and it is accompanied, for $T < T_c$, by a spontaneous breakdown of the flip-flop symmetry $\sigma_k^z \rightarrow -\sigma_k^z$ of the local Hamiltonians (104).
4. One ought to note at this point that the treatment in [Kac, 1968] was in the most orthodox spirit of classical statistical mechanics: an analysis by steepest descent methods of the partition function in the limit $|\Lambda| \rightarrow \infty$. The novelty in [Emch and Knops, 1970] was to consider also the evolution of the x - and y - components of the quantum spins and to study the resulting quantum dynamics in order to confront the interpretation of extremal KMS states with results already known from a classical treatment.

As with the BCS model reviewed in subsection 5.1 above, which has also the structure of a ‘molecular’ field model, some technicalities are involved here: in the thermodynamical limit, the convergence of the evolution is established only for the von Neumann algebras belonging to the representations considered.

5. This simplified version of the model, where (105) is assumed, already allows us to demonstrate the general features explored in this subsection. We may nevertheless mention that in [Emch and Knops, 1970] $J_{\Lambda, jk}$ is allowed to depend on the distance $|j - k|$, but only in such a way that for each $k \in \mathbb{Z}$, there exists a constant c_k such that $\sum_k |J_{\Lambda, jk}| < c_k$ for every finite Λ containing k , subject to the condition that $\forall j, k \in \mathbb{Z} : \lim_{|\Lambda| \rightarrow \infty} J_{\Lambda, jk} = 0$; the set of thermodynamical phases then becomes more complex, but its description still illustrates the adequacy of the decomposition account of spontaneous symmetry breakdown.
6. We already pointed out — first in section 4.1 — that as a consequence of the Liouville theorem of complex analysis, KMS states are necessarily time-invariant. Nevertheless, as established in [Emch and Knops, 1970], the present model admits *extremal KMS states that are not extremal time-invariant* — i.e. KMS states that are convex sums of different time-invariant, but not KMS, states — although these extremal KMS states do satisfy a very strong clustering property with respect to *space-translations*. The occurrence of extremal KMS states that are not extremal time-invariant reflects the fact that the time-evolution is *not* asymptotically abelian. This is not an unexpected peculiarity of the model since experience shows that asymptotic abelianness for the group of *time-translations* is rarely satisfied in QSP — although a few exceptions are known, among them the even part of the XY-model discussed at the beginning of subsection 5.4, despite the fact that

locality entails very strong asymptotic abelianness for the group of *space-translations*.

The coexistence of liquid and gas — say vapour and liquid water — presents formal similarities with the coexistence of magnetic phases oriented in opposite directions. The lattice-gas models of classical statistical mechanics are treated in close analogy with those of their ferromagnetic counterparts: instead of attributing to each site of a regular n -dimensional Ising model, a classical spin taking the values $+1/2$ and $-1/2$, one considers a random variable indexed by the sites of the lattice and taking the values 1 or 0 depending on whether the site is occupied by a molecule or not; double (or higher) occupancy is ruled out by *fiat* in these models. Phenomenologically, liquid-vapour coexistence curves in the phase diagram translate closely to the coexistence curves in ferromagnetic materials. In particular both present a critical point, precisely located in the phase space by the occurrence of diverging fluctuations. For temperatures higher than the critical temperature, any distinction between liquid and gas is untenable, and this state of matter is best described as a fluid.

B. QSP brought to bear on the Landau argument

The situation encountered with the coexistence of fluid and crystalline phases of the same substance — say water in its fluid phase and ice phase — is phenomenologically very different from the situation presented by a gas-liquid phase transition. Here, no critical point has been located: the fluid-solid coexistence curve extends indefinitely as pressure and density are increased. A heuristic argument for the non-existence of a critical point for the fluid-solid coexistence curve was advanced by Landau; see for instance [Landau and Lifshitz, 1958b, p. 260]. The argument was taken up by Uhlenbeck in [Uhlenbeck, 1968, p. 17]: “Because the solid and the fluid are with respect to long range order qualitatively different, there cannot be a critical point, since by going around it this would imply that long range order would appear gradually, which is impossible. This is the argument of Landau and I find it completely convincing.” And yet, Uhlenbeck warns on the same page that “one cannot escape the fact (intuitively evident, although not proved!) that there is already long range order in the solid phase itself.”

In an impressive sequence of papers, Kastler *et al.* [1967] rose to the challenge; the various assumptions of asymptotic abelianness, pervasive in these papers, was shown to be dispensable in the version worked out in [Emch *et al.*, 1970], which is followed here.

The programme is to classify the extremal KMS states that appear in the decomposition of a Euclidean invariant KMS state. Let \mathcal{A} be the C^* -algebra obtained as the C^* -inductive limit of local algebras $\mathcal{A}(\Lambda)$ over an absorbing net \mathcal{F} of finite regions $\Lambda \subset \mathbb{R}^3$ (here ‘finite’ means finite volume: $|\Lambda| < \infty$). Let $\alpha : t \in \mathbb{R} \mapsto \alpha_t \in \text{Aut}(\mathcal{A})$ describe an evolution; let $\nu : g \in \mathbb{E}^3 \mapsto \alpha_g \in \text{Aut}(\mathcal{A})$ describe the action of the Euclidean group \mathbb{E}^3 ; and let φ be a KMS state on \mathcal{A} with respect to the evolution α for the temperature β ; φ is assumed to be invariant under the action of the Euclidean group, i.e. $\forall g \in \mathbb{E}^3 : \varphi \circ \nu_g = \varphi$; this

condition is motivated by the phenomenological expectation that the underlying interactions are Euclidean invariant.

It is convenient to assume further that φ is strongly transitive with respect to the action of \mathbb{E}^3 in the sense that the following two conditions are satisfied.

1. For any two states ψ and ψ' appearing in the decomposition of φ in extremal KMS states, there exists at least one $g \in \mathbb{E}^3$ such that $\psi' = \psi \circ \nu_g$.
2. For one — and therefore all — state ψ appearing in the decomposition of φ into extremal KMS states, the isotropy subgroup $G_\psi := \{g \in \mathbb{E}^3 \mid \psi \circ \nu_g = \psi\}$ contains at least three non-coplanar translations.

Note that for any $g \in \mathbb{E}^3$ and any ψ appearing in the decomposition of φ into extremal KMS states, the state $\psi_g := \psi \circ \nu_g$ also appears there; and that $G_{\psi_g} = g^{-1}G_\psi g$. Hence, up to conjugacy, all elements appearing in the decomposition of φ have the same symmetry. This conjugacy class is denoted G^φ , and is referred to it as the *intrinsic symmetry* of φ . It is the part of the Euclidean symmetry of φ that is preserved when φ is decomposed into its extremal KMS components. Consequently, condition (1) is essentially one of convenience: if it were not satisfied, one would first have to separate the decomposed states in classes of conjugate elements, and carry out the analysis sketched below for each class separately. Condition (2) excludes pathological cases which one does not want to consider here. Mathematically, it strengthens condition (1) to ensure that the orbit of each extremal state under the translation group $\mathbb{R}^3 \subset \mathbb{E}^3$ in the space of all states on \mathcal{A} is closed.

It is then proven in [Emch *et al.*, 1970] that a Euclidean-invariant KMS state φ that satisfies the above conditions must necessarily belong to one of the following four classes.

The first class obtains when φ is already extremal KMS, i.e. its intrinsic symmetry is the group \mathbb{E}^3 itself. This case occurs *exactly when* one — and thus all — of the following *equivalent* conditions is satisfied:

1. φ is extremal \mathbb{R}^3 invariant, i.e. cannot be decomposed into a convex combination of states that are invariant under all translations in \mathbb{R}^3 .
2. The spectrum of the generator \mathbf{P} of the unitary representation of \mathbb{R}^3 canonically associated to φ by the GNS construction consists of exactly one eigenvalue, namely $\mathbf{k} = 0$, and this eigenvalue is non-degenerate.
3. φ is uniformly clustering in space, i.e. : for every $\epsilon > 0$ and $A \in \mathcal{A}$ there exists a finite region of space $\Lambda \subset \mathbb{R}^3$ such that

$$(107) \quad \forall B \in \mathcal{A}(\Lambda^c) : |\varphi(AB) - \varphi(A)\varphi(B)| \leq \epsilon \|B\|$$

where $\mathcal{A}(\Lambda^c) \subset \mathcal{A}$ is the C^* -inductive limit of the local algebras $\mathcal{A}(\Omega)$ with $\Omega \in \mathcal{F}$ and $\Omega \bowtie \Lambda$, (i.e. $\Omega \cap \Lambda = \emptyset$); see definition 27, scholium 29 and corollary 30 above.

In view of these properties, a state φ belonging to this class is interpreted as a *fluid phase*.

To describe the other three classes, namely the strongly transitive Euclidean invariant KMS states that do *not* describe fluids, let us focus now on the notion of the *intrinsic translational invariance* of φ . For any state ψ that appears in the decomposition of φ into extremal KMS states, let G_ψ denote the subgroup of Euclidean symmetries of ψ , and let $H_\psi = G_\psi \cap \mathbb{R}^3$ denote the subgroup of space-translations that preserve ψ . As one reviews the definition of the conjugacy classes one verifies that this group is indeed characteristic of the original state φ . Note also that strong transitivity entails that \mathbb{R}^3/H_ψ is compact.

The second class of Euclidean, strongly transitive KMS states is now specified by the following *equivalent* conditions, where ψ is any state appearing in the decomposition of φ into its extremal KMS components.

1. G_ψ is a crystallographic group.
2. φ is not extremal \mathbb{R}^3 -invariant, and H_ψ is generated by three non-coplanar translations.
3. With $\chi = \eta^{\mathbb{R}^3}[\psi]$ — where $\eta^{\mathbb{R}^3}$ is any invariant mean over the translation group \mathbb{R}^3 — χ is η -clustering (see definition 22 above), but neither weakly mixing nor even partially weakly mixing, i.e. χ satisfies

$$(108) \quad \forall A, B \in \mathcal{A} : \eta^{\mathbb{R}^3} (\chi(\nu_\bullet[A]B) - \chi(A)\chi(B)) = 0$$

but does *not* satisfy any of the stronger conditions

$$(109) \quad \forall A, B \in \mathcal{A} : \eta^{\mathbb{R}^3} |\chi(\nu_\bullet[A]B) - \chi(A)\chi(B)| = 0$$

$$(110) \quad \forall A, B \in \mathcal{A} : \eta^{\mathbb{R}^1} \left| \eta^{\mathbb{R}^2} (\chi(\nu_\bullet[A]B)) - \chi(A)\chi(B) \right| = 0$$

$$(111) \quad \forall A, B \in \mathcal{A} : \eta^{\mathbb{R}^2} \left| \eta^{\mathbb{R}^1} (\chi(\nu_\bullet[A]B)) - \chi(A)\chi(B) \right| = 0$$

REMARKS 64.

1. Taken separately, each of the conditions (1–3) excludes that φ be a fluid phase. Indeed, a fluid phase is extremal KMS, so that its intrinsic symmetry is the Euclidian group \mathbb{E}^3 , contradicting (1); a fluid phase is extremal \mathbb{R}^3 –invariant, contradicting (2); a fluid phase is uniformly clustering (see 5.33), which implies that each of the relations (108–111) would be satisfied, whereas (109–111) are not satisfied in the present phase.
2. The other two classes to which φ may belong are characterized as follows. (109), if satisfied, would have entailed $H_\psi = \mathbb{R}^3$, thus contradicting the second part of condition (2). This would correspond to a situation where the rotational symmetry is broken whereas the translational symmetry of the state φ would be completely preserved in its decomposition into its extremal KMS components. Although this may occur in systems exhibiting spontaneous magnetization, it is not immediately relevant to the purpose of identifying the way in which the formalism distinguishes fluids from solids in a world where fundamental interactions are invariant under the Euclidean group \mathbb{E}^3 .

Similarly, (110) or (111), if satisfied, would have entailed H_ψ is continuous in one or two direction(s) but discrete in the complementary direction(s). Such situations have been envisaged also — as early as the mid 1930s, cf. e.g. [Landau and Lifshitz, 1958b, p. 410] — but here again, their putative existence does not bear directly on the problem at hand.

3. The space-averaged state χ , which is \mathbb{R}^3 –invariant by construction, nevertheless keeps a memory of the symmetry of the state ψ from which it is constructed. Indeed the *discrete* part of the spectrum of the generator \mathbf{P}_χ of the unitary group representation of \mathbb{R}^3 , associated to the GNS construction corresponding to χ , coincides with the reciprocal group of H_ψ , namely with

$$(112) \quad H_\psi^* = \{\mathbf{k} \in \mathbb{R}^3 \mid \forall \mathbf{a} \in H_\psi : \mathbf{k} \cdot \mathbf{a} = 0 \pmod{2\pi}\}$$

which, in principle, is observable in X-ray diffraction patterns.

Upon keeping from the above what is relevant to the absence of a critical point in the coexistence curve between fluid and solid, the analysis of the decomposition of Euclidean invariant canonical equilibrium states into their pure thermodynamical phase components provides a rigid relation between the clustering properties and the geometric properties of these states. Namely: fluid phases exhibit a uniform clustering property (107), while crystalline phases lead to the distinctly weaker property of weak-clustering (108), thus vindicating the Landau argument.

6 WHENCE AND WHITHER QSP?

This final section may serve as a summary, a conclusion, a collection of appendices sharpening some aspects of the theory that have not been discussed in the main text; and hopefully, as a constructive prospectus for territories beyond the scope of this essay.

Let me summarize very briefly the story so far. First, recall that sections 1 to 3 reviewed some of the salient features that are variously treated in traditional texts on QSP. Then, in sections 4 and 5, I argued that the algebraic formalism of the KMS condition provides a well-defined syntax, the semantics of which supports the following associations in equilibrium QSP.

- Canonical equilibrium states are described by KMS states, a notion that translates naturally from finite systems to systems considered in the thermodynamical limit.
- Pure thermodynamical phases are described by extremal KMS states.
- Pure thermodynamical phase components that appear in the unique decomposition of a canonical equilibrium state may have a symmetry lower than that of the original state: only the manifold of the different contributing phases reflects the original symmetry; cf. [Liu and Emch, 2005] where we describe the “decompositional account” of spontaneous symmetry breakdown in the quantum theory of phase transitions.

Against this background, the material of the present section is presented in four subsections. I first review the mathematical concept of a limit and its physical interpretation as used in the main text. I then discuss again the notion of macroscopic observables, taking here a perspective that opens on the next subsection: the quantum measurement problem. Finally, I present some remarks — prospective and/or revisionist? — on the pursuit of constructive confrontations between mathematical and theoretical physicists in order that they better inform the wider arenas where philosophers of science operate.

6.1 *Four limiting procedures in QSP*

In dealing with the topics just reviewed, and as early as in sections 2 and 3, at least four different types of limits were encountered, alone or in concert.

1. the classical limit $\hbar \rightarrow 0$;
2. the high temperature limit $T \rightarrow \infty$;
3. the thermodynamic limit $|\Lambda| \rightarrow \infty$;
4. the van Hove limit $\{\lambda \rightarrow 0 \text{ and } t \rightarrow \infty\}$ with $\tau := \lambda^2 t$ remaining finite.

As the philosophical legitimacy of each of these four limits (or ‘limiting procedures’) has been variously questioned elsewhere, I ought to specify again — in the vernacular, i.e. without an explicit mention of the traditional $(\epsilon, \delta_\epsilon)$ — that the limits were consistently understood in this essay to be *controlled limits* in the sense of mathematics: you give me a tolerance, and I tell you the price; the smaller the tolerance, the higher the price; but however small the error you are willing to tolerate, there is a price under which you are guaranteed that the article will be within what you decided you are going to tolerate. Mathematical physics adds to this the requirement that the “price” be expressed in currencies recognized by the putative laboratory technician. Let us examine successively the above four limits from this perspective.

1. The classical limit.

The Planck constant is a *fundamental physical constant*: in cgs units $h \simeq 6.62 \times 10^{-27}$ ergsec; the familiar notation $\hbar := h/2\pi$ is used here. To say that it is small is a “value judgement”, reflecting the energy scale which you believe is relevant for the problem you wish to discuss. To illustrate the working of limiting processes, and their physical meaning, let us examine a specific example, the classical limit of a typically quantum phenomenon, the tunnel effect in which a particle of energy E does “slip through” a barrier of height $V_o > E$. This effect was discovered in 1928 independently by Gamow and by Gurney and Condon [Gamow, 1928; Gurney and Condon, 1928; Gurney and Condon, 1929] in their search for an explanation of alpha-particle emission from heavy nuclei. The Josephson junction — an oxide layer sandwiched between two superconductors — is a more recent manifestation of this quantum phenomenon; cf. e.g. [Josephson, 1982]. Let us consider here the simplest model, quantum tunnelling through a square one-dimensional barrier.

One verifies immediately that the Schrödinger equation

$$(113) \quad \left[-\frac{1}{2m} \hbar^2 \frac{d^2}{dx^2} + V(x) \right] \Psi(x) = E\Psi(x) \quad \text{with}$$

$$V(x) = \begin{cases} 0 & \text{if } x < a \\ V_o & \text{if } -a < x < a \\ 0 & \text{if } x > a \end{cases} \quad \text{where } 0 < a < \infty \text{ and } 0 < V_o < \infty$$

admits, whenever $0 < E < V_o$, a solution of the form

$$(114) \quad \Psi(x) = \begin{cases} A_- e^{ikx} + B_- e^{-ikx} & \text{if } x < a \\ A e^{\kappa x} + B e^{-\kappa x} & \text{if } -a < x < a \\ A_+ e^{ikx} & \text{if } x > a \end{cases}$$

with

$$(115) \quad k = \left\{ \frac{2mE}{\hbar^2} \right\}^{\frac{1}{2}} \quad \text{and} \quad \kappa = \left\{ \frac{2m(V_o - E)}{\hbar^2} \right\}^{\frac{1}{2}}$$

where the relative ratio of the five coefficients A_-, B_-, A, B, A_+ in (114) are determined by imposing four conditions, namely that Ψ and its derivative be continuous

at the boundaries $x = \pm a$. In particular, these conditions imply

$$A_- = A_+ e^{ika} \frac{1}{4ik\kappa} [(\kappa + ik)^2 e^{-2\kappa a} - (\kappa - ik)^2 e^{2\kappa a}] .$$

Then from the reverse triangle inequality $|a - b| \geq \max\{|a| - |b|, |b| - |a|\}$:

$$\left| \frac{A_-}{A_+} \right| \geq \frac{\kappa^2 + \kappa^2}{4k\kappa} (e^{2\kappa a} - e^{-2\kappa a}) = \left[\frac{1}{2} \frac{V_o}{\sqrt{E(V_o - E)}} \right] \sinh 2\kappa a$$

Since the term [...] is independent of \hbar , let us emphasize the role of \hbar by rewriting the above formula as:

$$(116) \quad \frac{|A_+|^2}{|A_-|^2} \leq C [\sinh 2\kappa a]^{-2} .$$

From the definition of A_{\pm} in (114), the left-hand side of (116) is to be interpreted as the transmission coefficient of the barrier. In the corresponding classical model, $0 < E < V_o$ entails that this coefficient vanishes. Thus, to demand that the quantum model approximates its classical counterpart is to require that the quantum transmission coefficient be arbitrarily small, say

$$(117) \quad \frac{|A_+|^2}{|A_-|^2} \leq C [\sinh 2K]^{-2}$$

with K as large as one desires. To ensure that (117) is satisfied, the computation leading to (116) shows that it is sufficient to have: $\kappa a > K$, i.e.

$$(118) \quad \hbar < K^{-1} [2m(V_o - E)]^{\frac{1}{2}} a .$$

Hence, the classical limit of the quantum system (113) now is controlled:

- (i) mathematically, through the conjunction of (117) and (118);
- (ii) physically, as (118) gives an estimate of its range of validity in terms of the physical quantities that characterize this system

In this sense the classical limit is similar to the non-relativistic limit: the classical description emerges from the quantum theory in the same way as Newton's mechanics emerges from Einstein's special relativity theory. The key to a proper understanding is the evaluation of the domain of validity of the approximations. Having done so, I have no qualm assuring my insurance agent that my car is not going to tunnel out of my garage, nor do I worry about relativistic red-shift when I drive my car in congested traffic. Such are the physical parameters that prevail in my car, my garage and the town where I live; compare with [Gamov, 1940] where Gamow pretends with didactic gusto that $\hbar \simeq 1$ ergsec or $c \simeq 15$ km/hour.

2. The high temperature limit.

Following the Ehrenfests, I have repeatedly argued in this essay that in QSP, the

classical regime emerges when the temperature is high enough. Typically, the quantities that tell us the regime in which we operate are similar to (βh) with $\beta = 1/kT$ (where k is the Boltzmann constant $k \simeq 1.38 \times 10^{-16}$ erg degree $^{-1}$).

To illustrate this point, let us review the results on the black-body radiation (subsection 2.1) and the specific heat of solids (subsection 2.3).

We saw qualitatively that if $h\nu \gg kT$ Planck's formula of 1900, here (1), reproduces (5) which had been previously established by Wien in 1896 [Wien, 1896]. Quantitatively, Paschen and Wanner [Paschen and Wanner, 1899] had verified in 1899 that Wien's formula is in agreement with laboratory data in the range of visible light, i.e. for wave length $\lambda = c/\nu$ between 4000   and 7000  , for temperatures up to 4000 K; this is what we would call today the "quantum regime". As the temperature at the surface of the sun is about 6000 K, going to much higher temperatures was then not an option. Nevertheless, as the ratio $h\nu/kT$ is concerned, raising T or lowering ν have the same effect; the latter means pushing the observation into the infrared, which was possible at the time. Indeed, the following year Lummer and Pringsheim [Lummer and Pringsheim, 1900] recorded systematic deviations from Wien's formula when the wavelength reaches the range of 12 to 18 μ (recall $1\mu = 10^{-6}$ m and thus $12\mu = 12 \cdot 10^4\text{  }$, compared with $\sim 7 \cdot 10^3\text{  }$ for visible red). This is the observation that prompted the purely classical derivation of the Rayleigh–Jones formula (6), and then Planck's interpolation between $h\nu \gg kT$ (Wien) and $h\nu \ll kT$ (Rayleigh–Jones). Experimentally, the passage from Paschen and Wanner to Lummer and Pringsheim thus marks very sharply in time (less than two years) the crossing of the boundary from the quantum to the classical regimes. These two regimes are numerically characterized by their distance $|\lambda - \lambda_{max}|$ from the wavelength λ_{max} — or equivalently the frequency ν_{max} — at which the Planck distribution (1) passes through a maximum.

As for the specific heat of solids, upon revisiting equations (8) and (9) Debye already verified that conclusion (10) can be sharpened to give the exact result

$$(119) \quad \text{where} \quad \left. \begin{aligned} C_V &= 3R\{4D(\frac{\Theta}{T}) - 3(\frac{\Theta}{T})[\exp(\frac{\Theta}{T}) - 1]^{-1}\} \\ D(x) &= \int_0^x dt \frac{t^3}{e^t - 1} \quad \text{and} \quad k\Theta = h\nu_o \end{aligned} \right\} .$$

Scholium 2 was obtained by noticing that

$$D(x) \simeq \begin{cases} 1 & \text{for } 0 < x \ll 1 \\ \frac{1}{5}\pi^4 x^{-3} & \text{for } x \gg 1 \end{cases} .$$

To go further than this, and determine the onset of the classical regime $C_V = 3R$, requires one to notice two things. First, C_V in (119) is a monotonically increasing *universal* function of the variable Θ/T ; while it cannot be written in terms of elementary functions, it can be computed numerically. Since Θ is known in term of the cut-off ν_o on the vibrational frequencies of the crystal, its value can be determined by mechanical means: for instance, at room temperature, Θ is about

100 K (for lead) and about 400 K (for aluminium), with silver and copper in between. For these, and many other metals, the measured values of the specific heat fall remarkably close to the theoretical prediction (119); cf. e.g. [Wannier, 1966, fig.13.9, p.276]. This curve shows a monotonic and smooth passage from the classical to the quantum regimes as the temperature decreases. Specifically, we can now discuss quantitatively the onset of the classical regime. The exact expression (119) entails that the first two terms in the expansion of C_V for $\Theta/T \ll 1$ give

$$(120) \quad C_V \simeq 3R\left\{1 - \frac{1}{20}\left(\frac{\Theta}{T}\right)^2\right\}$$

so that at room temperature $T \simeq 300$ K, the correction to the classical value $C_V = 3R$ ranges from about 0.6 % (for lead) to about 9 % for aluminium; both of which are in good agreement with experimental data.

For other early recognitions of the emergence of the classical regime in high temperature QSP, see subsections 2.4 and 2.6 where the classical ideal gas is recovered as controlled high temperature limits of both the Bose and the Fermi quantum gases.

3. The thermodynamical limit.

As its name indicates, the thermodynamical limit is designed to elicit various macroscopic thermodynamical behaviours from microscopic mechanical models. I find it convenient to separate here the problems addressed in non-equilibrium and in equilibrium statistical physics.

a. *Non-equilibrium physics.* In the classical realm already, one appeals to the large size of the systems considered to avoid the spurious appearance of recurrences in the theoretical modeling of physical phenomena such as the thermodynamical approach to equilibrium. For instance, to buttress Boltzmann's kinetic theory of gases, the Ehrenfests proposed the so-called dog-flea model, a stochastic model later revisited by Mark Kac. This model is reviewed in [Emch and Liu, 2002, section 3.4] where the results of a computer experiment are reported, involving $N = 100$ "fleas" jumping "at random" between two "dogs": a tendency to approach equilibrium is manifest during a few hundred jumps, whereas the frequency of recurrences, which Kac showed to increase exponentially with N , is observed to occur — as regularly as to be expected — over a range of several tens of thousands of trials.

In the quantum realm, a model for an actual experiment, the nuclear free-relaxation, is solved in subsection 3.3 above. Here again the model shows an approach to equilibrium practically unaffected by a "recurrence time" that grows as 2^N , where N is the number of lattice sites in the system, a macroscopic CaF_2 crystal. Hence the empirical justification for the limit $N \rightarrow \infty$ is that the relevant time-parameter for the experimentalist is $\sim 2^{10^{23}}$, which is indeed exorbitant; accordingly, I could not discern from the laboratory [Lowe and Nordberg, 1957] any concern about putative recurrences. The supporting analytic evidence is the explicit size-correction given in equation (49) and discussed with some detail in remarks 6.

b. *Equilibrium Physics*. In equilibrium situations, the thermodynamical limit is called upon to focus on properties of matter in bulk, so to speak navigating the high seas, away from the shoals of boundary effects. This often requires some elaborate rigging, in CSP as well as in QSP.

Roughly speaking, up to the middle of the twentieth century, this was achieved by replacing sums by integrals, as in equations (13–14). As in other parts of mathematical physics, this mathematical procedure is usually well under control, albeit physics sometimes requires unusual precautions, as shown explicitly in the caveat of equation (50).

Later on, especially in the modeling of phase transitions, when the emergence of collective behaviour turned out to be essential for the understanding of the phenomena at hand, and when existence questions were raised, more sophistication was demanded. In particular, for the limit where the size of the system is allowed to go to infinity, dimension enters the play; and then, in particular, the shapes of the regions considered must be such that the ratio of the surface to the volume goes to zero: cubes are fine; sponges are not. As indicated in the various models presented in section 5, it is possible to carry out such limiting procedures explicitly and successfully. The simplest examples are lattice systems, say spins on a lattice \mathbb{Z}^d . Examples of continuous systems are also given in this section 5; yet, in general, such systems, say on \mathbb{R}^d , require extra technical care to ensure spatial uniformity and to avoid bundling effects; hence, in the latter cases the theory is not always as fully controlled as one may desire; see nevertheless [Sewell, 2002]. Interactions with extremely long range may pose further problems with regard to: (a) the definition of the limiting state; and (b) the control of the limiting time evolution; such situations were met in subsections 5.1 and 5.7.A.

4. **The van Hove limit.**

We encountered particular instances of this limiting procedure in subsections 3.5 and 5.4. But a more general discussion was postponed to the present section.

In a brilliant transposition of a theme van Hove had heard played to justify the Born approximation in the discussion of long-time asymptotic behaviour in scattering theory, he proposed in [van Hove, 1955] a variation allowing him to characterize a regime where the time-scale of the irreversible *macroscopic* phenomena is emphasized over the time-scale of the underlying reversible, Hamiltonian *microscopic* dynamics. Van Hove’s original presentations were conducted for specific models by means of perturbation techniques, carried to all orders, in which he selected for summation the “most divergent diagrams.” At first, his virtuoso performances drew considerable scepticism; cf. e.g. [van Kampen, 1962]. The main problem was to isolate the conditions under which the essentials of what would become a theory may emerge from the contingent diagrammatics attached to the solution of particular models. Systematic mathematical treatments are now available to show how a joint long-time/weak-coupling limit may lead from a conservative unitary evolution to a contractive dynamical semi-group; cf. e.g. [Martin, 1979; Davies, 1976a].

In terms of the focus of this subsection, namely the control of limits that allow one to ensure that the system considered is operating in a desired regime, here exponential decay, one aims at proving a result of the following form; cf. e.g. [Martin and Emch, 1975, section 4].

There exist finite constants $\tau_o > 0$ and $C > 0$ such that for $0 \leq \lambda^2 t \leq \tau_o$:

$$(121) \quad \left| \lim_{|\Lambda| \rightarrow \infty} (\Phi, U_{-t}^o U_t \Psi)_\Lambda - (\Phi, \exp(-[\Gamma + i\Delta]\lambda^2 t)\Psi) \right| \leq \lambda C$$

where $U_{-t}^o U_t$ describes the evolution in the so-called interaction picture, with $U_t^o = \exp(-iH_o t)$, $U_t = \exp(-i[H_o + \lambda V]t)$; $H_o, H + \lambda V, \Delta$ are self-adjoint operators, and Γ is not only self-adjoint, but also positive so as to describe *decay* in the time range $0 \leq \tau := \lambda^2 t \leq \tau_o$. Hence the term “long-time/weak-coupling limit”: when the coupling constant λ is small enough [i.e. the RHS of 121 is small] the evolution is approximated by the contractive semi-group $S(\tau) := \exp(-[\Gamma + i\Delta]\tau)$ with $\tau = \lambda^2 t \in [0, \tau_o]$, provided the time t is sufficiently large in the scale measured by $t \simeq \tau/\lambda^2$.

The separation of the total Hamiltonian $H = H_o + \lambda V$ into an “unperturbed” or “free” part H_o and an “interaction” λV must be justified. Van Hove proposed that it is to be traced back to the fact that the observables of interest in irreversible processes are macroscopic (see subsection 6.2 below), thus determining a joint spectral resolution; then H_o appears as the “diagonal” part of H in this spectral resolution. For instance, $A = \int dk A(k) a^*(k) a(k)$ and $H_o = \int dk \epsilon(k) a^*(k) a(k)$. This remark also helps justify the use of the interaction picture $U_{-t}^o U_t$ since it entails that the observables of interest are invariant under the “free” evolution. As the macroscopic observables are translation invariant, the notation $\int dk$ is used to suggest that the momentum representation corresponds to the spectral resolution in which the observables and the free Hamiltonian are diagonal.

The understanding of the van Hove limit gained in the 1970s has since been confirmed and extended; cf. e.g. [Bach *et al.*, 2000; Dereziński and Früboes, 2005] and references therein; for baselines [Davies, 1976a], [Emch and Liu, 2002, section 15.2] and [Alicki and Fannes, 2001].

Although I do not wish to elaborate on the following historical point, I may mention incidentally that the use of the interaction picture helped van Hove discern in his perturbation expansions some *characteristic features* of many-body physics by which he suggested non-equilibrium QSP differ from the QFT supporting quantum scattering theory. To this day, however, I am not sure whether van Hove’s Delphian utterances have been properly digested into the corpus of contemporary mathematical physics.

I should also mention here that coupled limits have been considered also in CSP. An example is the Grad limit for classical gases in which the volume V is kept fixed, the number of molecules $N \rightarrow \infty$, and the cross-section of the molecules $\sigma := \pi d^2 \rightarrow 0$ (thus the volume of each molecule $v(\sim d^3) \rightarrow 0$ and the density of the gas $\rho := \frac{N}{V} \rightarrow \infty$), while the mean-free path $\lambda = \frac{V}{Nd^2}$ is kept constant; in [Grad, 1958], Grad proposed this limit as a mean to derive the Boltzmann

equation. For further references relative to the latter problem, see [Uffink, 2006, section I.6.2] or [Emch and Liu, 2002, section 3.3]; and in particular, for the 2-dimensional Lorentz gas (with now $\sigma = 2d$), cf. [Martin, 1979] where it is pointed out that the Grad limit and the van Hove limit (in a form of it adapted to this model) are equivalent in one important sense: they both predict the same ratio between the macroscopic time-scale validated by observations, and the microscopic time-scale provided by the mean free time between two successive collisions.

The related philosophical issues about the roles that asymptotic reasoning plays in explanation, reduction, and emergence are cogently discussed in [Batterman, 2002a]. The above four limiting procedures may bring additional water to this epistemological mill; see already [Grad, 1967].

In closing this subsection, I should at least mention *coarse-graining*, yet another procedure that has been transferred from the classical to the quantum realm [van Kampen, 1954; Emch, 1964]. One of the reasons for not bringing it up in this essay is that I did not need it for the considerations I developed here. And the reason for this may be that I have come to believe that the primacy of coarse-graining has been largely superseded by the syntax of infinite systems which allows one to bypass several awkward issues about the relations between the micro- and macroscopic worlds; see for instance in subsection 6.3 below. Yet, as with the thermodynamical limit, coarse-graining helped explore those macroscopic properties one wishes to see emerging from finer descriptions; in so doing, it also emphasizes that distinguishing differences of *scales* or *tempi* enables smooth negotiations of such passages.

6.2 Macroscopic observables

Coming back to the general formalism, assuming that the thermodynamic limit has been taken, and concentrating on space-translations, subsection 3.5 emphasized one feature that is new to *quantum* ergodic theory. Space-averaged observables are *essential observables* in the sense of the theory of superselection sectors prompted by [Wick *et al.*, 1952], i.e. they commute with all quasi-local observables and among themselves. This is yet another classical aspect of quantum theory. The specific classical description that emerges in this manner depends on the global preparation of the system (but is insensitive to local perturbations) as the very definition — and values — of these space-averaged observables depends on the translation-invariant state φ of the system one considers. This aspect of quantum ergodic theory shows up as a direct consequence of the “locality” assumed in the Haag–Kastler axioms.

Hence it is proper to regard the emergence of a classical macroscopic description out of a quantum microscopic description as a consequence of translation-invariance and locality; cf. subsections 3.5.B and 5.7. As we saw, the passage to the thermodynamical limit and the attendant emergence of macroscopic observables allow one to discern the simultaneous existence of several thermodynamical

pure phases, as for instance the non-vanishing magnetization in zero magnetic field signals the presence of a permanent magnet. Similarly, the laboratory observation of a discontinuity in the derivative of the isotherms at the ends of the Maxwell plateau is better understood if one takes the thermodynamical limit: otherwise, the isotherms are analytic all along and the theoretical description of their experiencing so extreme a bend is simply neither convenient nor useful when considering a cup of tea. And again, nobody would claim that when receiving their drinks they recognize the ice-cubes only because the size of these is infinite ... which, mercifully, it isn't. Yet, the Landau criterion for distinguishing a solid from a fluid (see subsection 5.7.b) is strictly valid only when the thermodynamical limit is considered. This is a paradox only when the definition of limits is forgotten; here as elsewhere in physics, the key to the proper understanding of limits lies in their manifesting the emergence of qualitatively different regimes.

For time-averaged observables, the situation is more complicated. Recall some basic facts. To any time-invariant state φ the GNS construction associates a representation π_φ of the C^* -algebra \mathcal{A} of quasi-local observables, and a unitary representation of the time evolution under which the von Neumann algebra obtained as the weak-operator closure of $\pi_\varphi(\mathcal{A})$, namely $\mathcal{N}_\varphi := \pi_\varphi(\mathcal{A})''$, is stable. Then, while the time-average of an observable always belongs to \mathcal{N}_φ as does its space-average, the time-average now also belongs to the commutant $\mathcal{N}_\varphi' = \pi_\varphi(\mathcal{A})'$ of this algebra, and thus to its center $\mathcal{Z}_\varphi := \pi_\varphi(\mathcal{A})'' \cap \pi_\varphi(\mathcal{A})'$, if and only if the evolution is η -abelian. The latter condition — see equation (61) — may be satisfied in some particular models, but its status is as yet too precarious to enshrine this condition as a general “axiom” on the same footing as “locality”.

In spite of the limitation just described, some of the remaining ergodic properties of observables under the time-evolution, together with some of the applications of the theory, were discussed in subsection 3.5.A.

6.3 *The quantum measurement process viewed from the perspective of QSP*

The technical literature on quantum measurement underwent some striking developments in the 1970s — cf. e.g. [Hepp, 1972; Bell, 1975; Whitten-Wolfe and Emch, 1976]; and also [Emch, 2003; Sewell, 2005] — in part as a consequence of the advent of the algebraic approach to QSP.

Insofar as there was a consensus on what the problem was, the original doctrine is best expounded in Wigner's careful exegesis of what he called the “orthodox” theory of von Neumann [von Neumann, 1932c]; Wigner's papers are collected in [Wigner, 1997, Part II] and [Wigner, 1995, Part II]; Wigner's positions on the subject were last stated in [Wigner, 1984]. Some of the philosophical issues are outlined in [Dickson, 2006].

A renewal in the understanding and implementation of several of the basic tenets of the doctrine was largely motivated by two critiques repeatedly advanced by Wigner himself. The first critique was that “to increase the accuracy of the

measurement one has to use a very large measuring apparatus” [Wigner, 1995, p. 177] or “the large size of the apparatus appears to be essential for the possibility of a measurement” [Wigner, 1995, p. 178]. The second critique is the problem of infinite regress — the so-called Wigner’s friend argument; cf. e.g. [Wigner, 1995, p. 215] — that follows from the necessity “to consider the system that has been called, so far, the apparatus, to be the object of the measurement. In other words, one will bring this apparatus into interaction with a new measuring object ... [*and so on*]”; [Wigner, 1995, pp. 208-9]. As this does not appear to be a problem with which one is usually concerned in the analysis of classical measurements, Wigner reiterated a statement he attributed to Fock, but which he said he believed to be part of the teaching of the “Copenhagen school”, namely that: “Measuring instruments must be described classically”; of singular relevance to the present essay, this quote is taken from a paragraph Wigner entitled “Is the measuring apparatus macroscopic?” [Wigner, 1995, p. 205].

The reason I believe to be at the core of this awkwardness is that in Wigner’s heydays, physicists were still in awe of a *perceived dichotomy* between the classical and the quantum worlds. Hence a new branch of the literature on the quantum measuring problem could develop when a solution of continuity was found that bridges these two descriptions — quantum and classical — of *the one world* in which we live. This happened when the conceptualization of the physical role of limiting procedures came under control and, in particular, the concept of macroscopic observables was understood; see subsections 6.1 and 6.2 above, references therein, and [Landsman, 2006]. I claim that the concepts developed to deal with QSP can help construct a measuring apparatus that is described in quantum terms and yet behaves, *qua* measuring apparatus, in a classical regime. I will now indicate how at least this part of the conceptual problems associated with quantum measurement has been clarified.

Let \mathcal{A}_S be the algebra of observables for the system to be measured, and let $\mathcal{B} \subset \mathcal{A}_S$ be an abelian subalgebra, the self-adjoint elements of which are the observables of interest. In the interest of formal simplicity I make here the following assumptions, parts of which are easy to dispense with.

- \mathcal{A}_S contains a unit I_S and is a collection of finite-dimensional matrices.
- The spectrum of \mathcal{B} is non-degenerate; hence every observable $B \in \mathcal{B}$ is of the form $B = \sum_k b_k Q_k$ with $Q_k = Q_k^*$; $Q_k Q_l = \delta_{kl} Q_k$; $\sum_k Q_k = I$; and $\dim Q_k = 1$.

Initially, the system of interest is in the state $\varphi_S : A \in \mathcal{A}_S \mapsto \text{Tr} \rho A \in \mathbb{C}$, and we want the measuring process to determine, for all $B \in \mathcal{B}$ the values $\varphi_S(B)$, i.e. for all k , the values $\lambda_k = \varphi_S(Q_k)$, so that we can compute $\varphi_S(B) = \sum_k b_k \lambda_k$.

For this measurement, a team of quantum engineers will be asked to build a dedicated measuring apparatus described by an algebra \mathcal{A}_M with self-adjoint “pointers” M_k which are in bijective correspondence with the Q_k . They prepare

this apparatus in the state φ_M . For simplicity, they assume that their \mathcal{A}_M contains a unit I_M and that they arrange for $\sum_k M_k = I_M$. And finally, they try to build an interactive *Hamiltonian* mechanism such that when the system of interest and the apparatus are brought into contact the initial state $\varphi^o = \varphi_S \otimes \varphi_M$ on $\mathcal{A}_S \otimes \mathcal{A}_M$ will evolve in such a manner that the following two conditions are satisfied:

(a) concerning the measuring apparatus:

$$(122) \quad \forall M_l : \begin{cases} \varphi_M(M_l) \longrightarrow \varphi^p(M_l) = \sum_k \lambda_k \psi_k(M_l) & \text{where} \\ \psi_k(M_l) = \delta_{kl} & \text{with no dispersion} \end{cases} ;$$

(b) concerning the system to be measured:

$$(123) \quad \forall A_S \in \mathcal{A}_S : \begin{cases} \varphi_S(A_S) \longrightarrow \varphi^p(A_S) = \sum_k \lambda_k \varphi_k(A_S) & \text{where} \\ \varphi_k(A_S) = \begin{cases} \varphi_S(Q_k)^{-1} \varphi_S(Q_k A_S Q_k) & \text{when } \lambda_k \neq 0 \\ \varphi_S(A_S) & \text{when } \lambda_k = 0 \end{cases} \end{cases} .$$

Let me comment on these design requirements. Note first that (122) would deliver the values $\lambda_k = \varphi_S(Q_k)$ from which one computes the expectation values $\varphi_S(B)$ of all observables for the measurement of which the apparatus was designed. I will specify later — see (125) — what is meant by the requirement that the result of the measurement be “without dispersion”, i.e. formally $\varphi^p([X - \varphi^p(X)I]^2) = 0$.

To relate the requirement (123) to the familiar textbook description of the measuring process, consider briefly the particular form it takes in the von Neumann framework where φ_S is a pure state on the algebra $\mathcal{A}_S = \mathcal{B}(\mathcal{H}_S)$ and the Q_k are one-dimensional; let $\{\Phi_k\}$ be an orthonormal basis in \mathcal{H}_S with $Q_k \Phi_l = \delta_{kl} \Phi_k$; in term of this basis, one can write, without loss of generality $\varphi_S(A_S) = (\Phi_S, A_S \Phi_S)$ with $\Phi_S = \sum_k c_k \Phi_k$; and $\lambda_k = |c_k|^2$. Then (123) takes the form $\varphi^p(A_S \otimes I_M) = \text{Tr}(\rho^p A_S)$ with $\rho^p = \sum_k |c_k|^2 Q_k$. Hence, viewed from \mathcal{A}_S , the pure state-vector Φ_S evolves to the mixed density matrix ρ^p . In this sense, (123) is the general form of the so-called von Neumann (non-selective) collapse postulate for the case where the initial state of the system is not necessarily a pure state.

Note that (122) and (123) are reduced descriptions of the evolution of the state φ^o : these requirements demand only that the evolution of special observables be followed; these special observables are: (a) the pointers M_l of the apparatus; and (b) all observables A_S pertaining to the system S . In particular, the requirement (123) would not be incompatible with a measuring process (which we denote as \longrightarrow) driven by a unitary evolution of the composite *system* \cup *apparatus*.

In line with von Neumann’s “relative frequencies” view of quantum probability — explicitly inspired by von Mises [von Neumann, 1932c, fn. 156] — the general form (123) applies best to a measurement performed on a beam of particles rather than separately on individual particles. Hence — in line with the interpretation of ‘states of physical systems’ stated in subsection 3.1 — this description of the measuring process understands that the initial state of the system S is viewed

as a summary of its preparation. For instance, in the historical Stern–Gerlach experiment, an incident *beam* of silver atoms was produced by evaporation from a heated oven; cf. [Jammer, 1966, p. 133]. Thus, what the experimentalists knew was the direction of the beam and the temperature of the oven: the latter surely a macroscopic notion! Similarly, the initial state of the measuring apparatus is viewed here as the result of its preparation; adhering to this pragmatic interpretation, one ought not to impose on the initial state of a (large!) measuring apparatus that it be pure: plainly this would require an exorbitant amount of information to be entered in its preparation — information that ought not to be actually necessary for the adequate performance of measurements aiming to collect the simple microscopic information described by the distribution $\{\lambda_k\}$.

Due to all sorts of pesky circumstances — e.g. the recurrences present in finite systems or the intrusion of the “Wigner’s friend” (introduced earlier in this subsection) — our apparatus builders would be exposed to dire frustrations, unless they be granted enough time and space so that the following idealization is a close enough approximation — to a degree chosen in advance — of their implementation of the measuring process $\varphi^o \longrightarrow \varphi^p$, namely:

$$(124) \quad \varphi^p(X) := \lim_{t \rightarrow \infty} \lim_{|\Lambda| \rightarrow \infty} \varphi^o(\alpha_t^\Lambda[X]) \quad \text{with} \quad X = \begin{cases} A_S \otimes I_M \\ \text{or} \\ I_S \otimes M_k^\Lambda \end{cases},$$

choosing the pointers so that in the thermodynamical limit, $\lim_{|\Lambda| \rightarrow \infty} M_k$ exist and define ‘essential’ observables — in the sense of subsection 6.2 above; in particular, the reader may want to review the connection with superselection rules — i.e. observables that the orthodox theory would construe to be classical. The requirement “without dispersion” in (122) may now be specified, namely one demands that

$$(125) \quad \lim_{t \rightarrow \infty} \lim_{|\Lambda| \rightarrow \infty} [\varphi^o(\alpha_t[(M^k)^2]_\Lambda) - \{\varphi^o(\alpha_t[(M^k)]_\Lambda)\}^2] = 0 \quad .$$

There is even an additional benefit in allowing the thermodynamical limit in (124), namely that one may demand that the experimental set-up be such that the result (124) of the measurement be empirically insensitive to local perturbations in the preparation of the initial state φ_M of the apparatus. This requirement means that $\varphi^p(X)$ in (124) do not change when the initial state φ_M of the apparatus is replaced by any state $\psi_M : A \in \mathcal{A}_M \mapsto \varphi_M(D^*AD) \in \mathbb{C}$ where D is any (quasi-)local element of \mathcal{A}_M satisfying the normalization $\varphi(D^*D) = 1$; or, even more generally, by any state ψ_M normal on the von Neumann algebra $\pi_{\varphi_M}(\mathcal{A}_M)''$. Such robustness pertains to the pragmatic demand that the preparation of a large(!) measuring apparatus be reasonably simple.

Here ends — at least for the main purpose of this subsection — the list of specifications demanded from our quantum engineers when constructing a measuring device.

The contribution of algebraic QSP to the solution of the quantum measurement problem is this: the above programme can be completely implemented in the

sense that specific and rigorously controllable models have been built satisfying *all* of the above specifications. These models therefore establish the applicability of the algebraic approach to the foundations of physics beyond the limitations of what Wigner called the orthodox theory. In sum, this approach encompasses the description of classical regimes unknown within the confines of the orthodox theory; cf. e.g. [Hepp, 1972; Whitten-Wolfe and Emch, 1976; Emch, 2003; Sewell, 2005] and other references listed in [Landsman, 2006, subsection 6.6].

An objection to (124), namely that real-world laboratories are finitely extended in space and in time is seductive. But it neglects the main understanding that presides over taking a limit: recall subsection 6.1 above. Here also the limit defines an asymptotic regime; thus, the control of the limiting procedures allows to take into account that good experiments do require expenses in room and allotments of time, each to be evaluated in terms of the precision to which one aims. The measuring process involves a particular instance of a general macroscopic phenomenon, the “approach to equilibrium”. In subsection 6.1(3) above, I commented again on the role of the thermodynamical limit $|\Lambda| \rightarrow \infty$ in the emergence of this regime.

The role of the subsequent limit $t \rightarrow \infty$ deserves a further comment in the context of the measurement process: it does *not* say that an infinite time is required to register the result of the measurement, but rather, in accordance with our general understanding of the role of limits, the existence of the limit $t \rightarrow \infty$ asserts that for every $\epsilon > 0$, there exists a time T_ϵ that can be evaluated, and is such that the measurement has been completed *for ever, within the required precision ϵ* , when $t > T_\epsilon$. Thus in contrast with the constraint of the orthodox theory requiring that the unitary evolution be sharply interrupted at the ‘end’ of the measurement process, our quantum engineers do not need to make provisions for switching off the measuring device. Now, not taking first the limit $|\Lambda| \rightarrow \infty$, only requires them to review their estimate of the effects of the finite size of the apparatus; from this estimate, they evaluate how large the apparatus must be so as to allow a generous time T_Λ before which they have to switch off the measurement and avoid some nasty kickback. The *controlled* limit $|\Lambda| \rightarrow \infty$ is thus not a pragmatic limitation to the validity of the theory any more than is the theoretical implementation of the thermodynamical limit ($N \rightarrow \infty$, $|\Lambda| \rightarrow \infty$ with $D := N/|\Lambda|$ fixed) to remove astronomically long recurrences from the description of the cooling down of your everyday cup of coffee. The description obtained in the thermodynamical limit is closer to the pragmatic account of the observed cooling down than would be its description as occurring in a finite system: the latter description would indeed be hampered by superfluous, irrelevant details. To sum up, in the actual construction of models for the measuring process, the problems that our quantum engineers encountered were not with satisfying the ancillary condition $T_\epsilon \ll \tau \ll T_\Lambda$ where τ denoted the laboratory time-scale on their wristwatch. See nevertheless [Bell, 1975].

While the models *do prove* that all the demands of the above programme are compatible, it is in the very nature of models that they *cannot* prove that

- (i) the conditions of the programme are necessary to an understanding of the measuring process; nor
- (ii) the conditions of the programme are sufficient, as other demands may be made, and other conditions may need to be required.

Concerning remark (i), the programme presented above emphasizes possible contributions that QSP can bring to an understanding of the quantum measurement process. One specific aim was to avoid having the theories of the measurement process beached on a conceptual sandbar between the quantum and classical worlds: the programme exploits circumstances where QSP shows how the quantum description of the one world encompasses conceptually important classical aspects. Thus the irreducible quantum/classical dichotomy has now faded into more comprehensive views, QSP being one of them. The emergence of classical behaviour in quantum theory is also one of the significant aspects of the *decoherence* programme, although the likely confluence of these two approaches has not yet gained universal acceptance. For a fair description of the latter issues, and their bearing on the measurement problem, I would recommend [Landsman, 2006]; and for a vivid and somewhat confrontational exchange on the relevance of decoherence in this context, [Anderson, 1994; Adler, 2003].

Remark (ii) above has at least two aspects. One of these aspects is that while the models that establish the internal consistency of the programme discussed in this subsection are treated with mathematical rigour, they can hardly be viewed as sufficiently realistic to satisfy our colleagues on the laboratory floor. Another aspect of the above remark (ii) on sufficiency, is that I do not know how the algebraic QSP would be helpful for formulating some of the remaining challenging questions still open in the theory of quantum measurement. If I had to single out one among these, I would direct attention first to measurements now “routinely” performed on an individual quantum system; cf. e.g. [Rauch and Werner, 2000] or [Rauch, 2005]. Whether the so-called “many worlds” and “consistent histories” approaches are really called for here is too wide a question to be addressed in this essay on QSP; cf. [Dickson, 2006; Landsman, 2006].

6.4 *Mathematical physics vs. theoretical physics*

Several largely unsolved problems may have been overshadowed by the abundant literature on the “return to equilibrium” of small or local deviations that are driven back to equilibrium by a thermal bath; for models of such coupled systems, see paragraphs A and C in subsection 6.4.

Most of the problems discussed below occur also in classical statistical physics; QSP offers little to alleviate them, but a little it does do, and here is how.

The first of these problems is to avoid an infinite regress: if a (small) system of interest is driven to equilibrium by a (large) thermal reservoir, whence is the reservoir getting its own canonical equilibrium and temperature? Rather than a

conceptual answer to this question, the KMS condition was originally conceived as a clever, but formal, transcription — from theoretical to mathematical physics — that turned out to be a wonderfully useful organizing tool.

This very success demanded that the KMS condition be given a deeper physical justification. Substantial answers were found later, diversely expressed as several stability conditions. The latter were presented in subsection 5.4 in an order in which their formulations increasingly sound more like *bona fide* variational principles. This development is thus in line with the widely held opinion that “a variational principle is considered to be the supreme form of a law of physics” [Itô, 1987, Art. 441]. This is good, but as in other fields, a philosophical question persists as to whether any science ought to be solely, or ultimately, founded on variational principles as mechanics and so many sciences have since the eighteenth century. Theoretical physics may have offered some other considerations in this regard, such as the “big-bang” and “decoherence,” but their explanatory value, consistency and adequacy remain to be proven. In the meantime, it is not unreasonable to prefer the updated variational principles with which algebraic QSP has proven able to refine their more traditional versions.

A second problem raised by the physics literature on the return to equilibrium concerns the description of global transport phenomena such as heat conduction and electric resistivity due to the interactions between electrons and phonons or random impurities in metals. Van Hove proposed a programme — of which the van Hove limiting procedure is a part, see subsection 6.1(4) above — to approach this type of question. One of the remaining problems is to produce mathematically clean arguments for the claims that are made. An even larger problem still to be fully mastered is to go beyond the contingencies of particular *ad hoc* models. This will require one to explain in physical terms amenable to a mathematical description the general microscopic properties actually responsible for a realistic delineation of the time-scales and/or regimes in which one observes such macroscopic phenomena; the first examples that come to mind are Newton’s “cooling law” and Fourier’s “heat” theory, i.e. the exponential temperature equilibration of temperatures and the flow of heat that governs the steady temperature distribution in materials placed between sources at different temperatures. The materials presented in this essay, particularly in subsections 3.5 [e.g. eqn. (57)] or 5.4 [e.g. eqn. (86)], exemplify some of the first steps that have been taken profitably along this road. Further, and promising but still formal, results have been obtained in [Eckman *et al.*, 1999; Bonetto *et al.*, 2000; Bach *et al.*, 2000], yet much remains to be done to bridge these with earthly concerns for an understanding that would allow one to compute realistic estimates of the value of specific material transport coefficients.

A third and perhaps more troubling problem. Time-reversal or not [Earman, 2002; Fredenhagen, 2003], even in my dreams I have not yet seen any “cosmological arrow of time” flying convincingly through the landscape of the C^* -algebraic approach developed for QSP ... but neither may such a flight be ruled out as a

heretical foray into this formalism [Buchholz, 2003].

A fourth direction in which to look for extensions of the programme of QSP is concerned with situations arising far away from equilibrium.

Yet a fifth arena for investigations has opened, where a connection with the algebraic approach to QSP is emerging. It will indeed be interesting to observe whether and how the maturing mathematical theory of quantum stochastic processes [Parthasarathy, 1995; Hudson, 1998] will or may throw new light on the reduction process of statistical mechanics.

Finally, QSP has of course found most of its pragmatic confirmation in the praxis of condensed matter physics and the extension of the latter into the study of complex phenomena. However, getting enmeshed here into the technical concrete details indispensable to the full mastery of this praxis would have carried us much beyond the confines of this essay. A richly documented overview of the scope of this field of enquiry may be found in [Anderson, 1994]. Yet, as with [Feynman, 1998], such matters need to be taken up again to weave in more threads and knots as well as to incite new philosophical reflections:

Vingt fois sur le m tier remettez votre ouvrage ...

BIBLIOGRAPHY

- [Adler, 2003] S. Adler. Why decoherence has not solved the measurement problem: a response to P.W. Anderson. *Studies in History and Philosophy of Modern Physics*, 34:135–142, 2003. See already [Anderson, 2001].
- [Aichele *et al.*, 2005] T. Aichele, U. Herzog, M. Scholz and O. Benson. *Single-photon generation and simultaneous observation of wave and particle properties*, pages 35–41 in [Khrenninkov, 2005].
- [Alicki *et al.*, 1998] R. Alicki, M. Bozejko and W.A. Majewski, editors. *Quantum probability*. Banach Center Publications 43, 1998.
- [Alicki and Fannes, 1994] R. Alicki and M. Fannes. Defining quantum dynamical entropy. *Letters in Mathematical Physics*, 32:75–82, 1994.
- [Alicki and Fannes, 2001] R. Alicki and M. Fannes. *Quantum dynamical systems*. Oxford University Press, Oxford, New York, 2001.
- [Amrein *et al.*, 1977] W.O. Amrein, J.M. Jauch and K.B. Sinha. *Scattering theory in quantum mechanics*. Benjamin, Reading MA, 1977.
- [Anderson, 1994] P.W. Anderson. *A career in theoretical physics*. World Scientific, Singapore, 1994.
- [Anderson, 2001] P.W. Anderson. Science: a ‘dappled world’ or a ‘seamless web’? *Studies in History and Philosophy of Modern Physics*, 32:487–494, 2001.
- [Andries *et al.*, 1995] J. Andries, M. Fannes, P. Tuyls and R. Alicki. The dynamical entropy of the quantum Arnold cat map. *Letters in Mathematical Physics*, 35:375–383, 1995. See already [Alicki and Fannes, 1992].
- [Arafune *et al.*, 2003] J. Arafune *et al.*, editors. *A garden of quanta, essays in honor of Hiroshi Ezawa*. World Scientific, Singapore, 2003.
- [Araki, 1961/2] H. Araki. Einf rung in die Axiomatische Quantenfeldtheorie, unpublished lecture notes at the ETH–Zurich, 1961/2.
- [Araki, 1969] H. Araki. Gibbs state of a one dimensional quantum lattice. *Communications in Mathematical Physics*, 14:120–157, 1969. See also [Araki, 1970] and [Araki, 1975].

- [Araki, 1970] H. Araki. *One dimensional quantum lattice system*, pages 75–86 in [Michel and Ruelle, 1970].
- [Araki, 1974] H. Araki. On the equivalence of the KMS condition and the variational principle for quantum lattice systems. *Communications in Mathematical Physics*, 38:1–10, 1974.
- [Araki, 1975] H. Araki. On the uniqueness of KMS states of one-dimensional quantum lattice systems. *Communications in Mathematical Physics*, 44:1–7, 1975. See also [Kishimoto, 1976].
- [Araki, 1984] H. Araki. On the XY-model on two-sided infinite chain. *Publications of the Research Institute for Mathematical Sciences, Kyoto University*, 20:277–296, 1984. See also [Araki and Barouch, 1983].
- [Araki and Barouch, 1982] H. Araki and E. Barouch. On the dynamics and ergodic properties of the XY-model. *Journal of Statistical Physics*, 31:327–345, 1983. See also [Araki, 1984; Araki and Matsui, 1985; Narnhofer, 1998].
- [Araki and Matsui, 1985] H. Araki and T. Matsui. Ground states of the XY-model. *Communications in Mathematical Physics*, 101:213–245, 1985.
- [Araki and Miyata, 1968] H. Araki and H. Miyata. On the KMS boundary condition. *Publications of the Research Institute for Mathematical Sciences, Kyoto*, A4:373–385, 1968.
- [Araki and Moriya, 2002] H. Araki and H. Moriya. Local thermodynamical stability of Fermion lattice systems. *Letters in Mathematical Physics*, 62:33–45, 2002.
- [Araki and Sewell, 1977] H. Araki and G.L. Sewell. KMS conditions and local thermodynamical stability of quantum lattice systems. *Communications in Mathematical Physics*, 52:103–109, 1977.
- [Araki and Woods, 1963] H. Araki and E.J. Woods. Representations of the canonical commutation relations describing a nonrelativistic infinite free Bose gas. *Journal of Mathematical Physics*, 4:637–662, 1963.
- [Arnold and Avez, 1968] V.I. Arnold and A. Avez. *Ergodic problems in classical mechanics*. Benjamin, New York, 1968.
- [Bach et al., 2000] V. Bach, J. Fröhlich and I.M. Segal. Return to equilibrium. *Journal of Mathematical Physics*, 41:3985–4060, 2000.
- [Bak, 1967] T.A. Bak, editor. *Statistical mechanics, foundations and applications; Proceedings of the IUPAP meeting, Copenhagen, 1966*. Benjamin, New York, Amsterdam, 1967.
- [Bardeen et al., 1957] J. Bardeen, L.N. Cooper and J.R. Schrieffer. Theory of superconductivity. *Physical Review*, 108:1175–1204, 1957.
- [Barton, 1963] G. Barton. *Introduction to advanced field theory*. Wiley, New York, 1963.
- [Batterman, 2002a] R.W. Batterman. *The devil is in the details: asymptotic reasoning in explanation, reduction, and emergence*. Oxford University Press, Oxford, New York, 2002.
- [Batterman, 2002b] R.W. Batterman. Asymptotics and the role of minimal models. *British Journal for the Philosophy of Science*, 53:21–38, 2002.
- [Bell, 1975] J.S. Bell. On wave packet reduction in the Coleman–Hepp model. *Helvetica Physica Acta*, 48:93–98, 1975.
- [Beller, 1999] M. Beller. *Quantum dialogue: the making of a revolution*. University of Chicago Press, Chicago, 1999.
- [Belot and Earman, 1997] G. Belot and J. Earman. Chaos out of order: quantum mechanics, the correspondence principle and chaos. *Studies in History and Philosophy of Modern Physics*, 28B:147–182, 1997.
- [Benatti et al., 1991a] F. Benatti, H. Narnhofer and G.L. Sewell. A non-commutative version of the Arnold cat map. *Letters in Mathematical Physics*, 21:157–172, 1991. Erratum in [Benatti et al., 1991b].
- [Benatti et al., 1991b] F. Benatti, H. Narnhofer and G.L. Sewell. Erratum: “A non-commutative version of the Arnold cat map”. *Letters in Mathematical Physics*, 22:81, 1991.
- [Besicovitch, 1954] A.S. Besicovitch. *Almost periodic functions*. Dover, New York, 1954.
- [Birkhoff, 1931] G.D. Birkhoff. Proof of the ergodic theorem. *Proceedings of the National Academy of Sciences (USA)*, 17:656–660, 1931.
- [Birkhoff, 1942] G.D. Birkhoff. What is the ergodic theorem? *American Mathematical Monthly*, 49:222–236, 1942. See also [von Neumann, 1932a].
- [Bisognano and Wichmann, 1975] J.J. Bisognano and E.H. Wichmann. On the duality condition for a Hermitian scalar field. *Journal of Mathematical Physics*, 16:985–1007, 1975. See also [Bisognano and Wichmann, 1976].
- [Bisognano and Wichmann, 1976] J.J. Bisognano and E.H. Wichmann. On the duality condition for quantum fields. *Journal of Mathematical Physics*, 17:303–321, 1976.

- [Boltzmann, 1884] L. Boltzmann. Ableitung der Stefanschen Gesetzes betreffend die Abh ngigkeit der W rmestrahlung von Temperatur aus der electromagnetischen Lichttheorie. *Annalen der Physik und Chemie*, 22:291–294, 1884.
- [Bonch-Bruevich and Tyablikov, 1962] V.L. Bonch-Bruevich and S.V. Tyablikov. *The Green function method in statistical mechanics*. North-Holland, Amsterdam, 1962.
- [Bonetto et al., 2000] F. Bonetto, J.L. Lebowitz and L. Rey-Bellet. *Fourier’s law: a challenge to theorists*, pages 128–150 in [Fokas et al., 2001].
- [Borchers, 1993] H.J. Borchers. On modular inclusion and spectrum condition. *Letters in Mathematical Physics*, 27:311–323, 1993.
- [Borchers, 1995] H.J. Borchers. On the use of modular groups in quantum field theory. *Annales de l’Institut Henri Poincar *, 67:331–382, 1995.
- [Borchers, 1999] H.J. Borchers. On the embedding of von Neumann subalgebra. *Communications in Mathematical Physics*, 205:69–79, 1999. See already [Borchers, 1993].
- [Borchers, 2000] H.J. Borchers. On revolutionizing quantum field theory with Tomita’s modular theory. *Journal of Mathematical Physics*, 41:3604–3673, 2000. See already [Borchers, 1995].
- [Bose, 1924] S.N. Bose. Plancks Gesetz und Lichtquantenhypothese. *Zeitschrift f r Physik*, 26:178–181, 1924.
- [Brout, 1965] R.H. Brout. *Phase transitions*. Benjamin, New York, Amsterdam, 1965.
- [Browder, 1976] F.E. Browder, editor. *Mathematical developments arising from Hilbert’s problems*, volume 28 in *Proceedings of Symposia in Pure Mathematics*, 1976.
- [Brunetti et al., 2003] R. Brunetti, K. Fredenhagen and R. Verch. The generally covariant locality principle — a new paradigm for local quantum field theory. *Communications in Mathematical Physics*, 237:31–68, 2003. See also [Fredenhagen, 2003].
- [Buchholz, 2003] D. Buchholz. On hot bangs and the arrow of time in relativistic quantum field theory. *Communications in Mathematical Physics*, 237:271–288, 2003.
- [Buchholz and Lechner, 2004] D. Buchholz and G. Lechner. Modular nuclearity and localization. *Annales de l’Institut Henri Poincar *, 5:1065–1080, 2004.
- [Buchholz et al., 2002] D. Buchholz, I. Ojima and H. Roos. Thermodynamic properties of non-equilibrium states in quantum field theory. *Annals of Physics*, 297:219–242, 2002.
- [Bunge, 1967] M. Bunge, editor. *Delaware seminar in the foundations of physics*. Springer, New York, 1967.
- [Caldirola, 1961] P. Caldirola, editor. *Ergodic theories*. Academic Press, New York, 1961.
- [Catto et al., 1998] I. Catto, C. Le Bris and P.-L. Lions, *The mathematical theory of thermodynamical limits: Thomas–Fermi type models*. Oxford University Press, Oxford, New York, 1998.
- [Chandrasekhar, 1931a] S. Chandrasekhar. The maximum mass of ideal white dwarfs. *The Astrophysical Journal*, 74:81–82, 1931. See also [Chandrasekhar, 1931b], [Chandrasekhar, 1931c], [Chandrasekhar, 1932].
- [Chandrasekhar, 1931b] S. Chandrasekar. The highly collapsed configuration of a stellar mass. *Monthly Notices Royal Astronomical Society*, 91:456–466, 1931.
- [Chandrasekhar, 1931c] S. Chandrasekar. The density of white dwarf stars. *Philosophical Magazine*, 11:592–596, 1931.
- [Chandrasekhar, 1932] S. Chandrasekar. Some remarks on the state of matter in the interior of stars. *Zeitschrift f r Astrophysik*, 5:321–327, 1932.
- [Chandrasekhar, 1958] S. Chandrasekhar. *An Introduction to the study of stellar structure*. Dover, New York, 1958. (1st ed. University of Chicago Press, 1939).
- [Chr tien et al., 1968] M. Chr tien, E.P. Gross and S. Deser, editors. *Statistical Physics, Brandeis Summer Institute, 1966*. Gordon and Breach, New York, 1968.
- [Churchill and Brown, 1990] R.V. Churchill and J.W. Brown. *Complex variables and applications*, 5th ed. McGraw-Hill, New York, 1990.
- [Clifton and Halvorson, 2001] R. Clifton and H. Halvorson. Are Rindler quanta real? Inequivalent particle concepts in quantum field theory. *British Journal for the Philosophy of Science*, 52:417–470, 2001.
- [Cohen, 1962/1968] E.G.D. Cohen, editor. *Fundamental problems in statistical mechanics, i, ii*. North-Holland, Amsterdam, 1962, 1968.
- [Connes, 2000] A. Connes. A short survey of noncommutative geometry. *Journal of Mathematical Physics*, 41:3832–3891, 2000.
- [Cornfeld et al., 1982] I. Cornfeld, S.V. Fomin and Ya.G. Sinai. *Ergodic theory*, Springer, Berlin, 1982.

- [Davies, 1972] E.B. Davies. Diffusion for weakly coupled oscillators. *Communications in Mathematical Physics*, 27:309–325, 1972.
- [Davies, 1974] E.B. Davies. Markovian master equations I. *Communications in Mathematical Physics*, 39:91–110, 1974.
- [Davies, 1976a] E.B. Davies. *Quantum theory of open systems*, Academic Press, London, 1976. See also [Davies, 1976b], [Davies, 1976c].
- [Davies, 1976b] E.B. Davies. Markovian master equations II. *Mathematische Annalen*, 219:147–158, 1976.
- [Davies, 1976c] E.B. Davies. Markovian master equations III. *Annales de l'Institut Henri Poincaré*, 11:265–273, 1976. See also [Gorini *et al.*, 1978].
- [Davies, 1975] P.C.W. Davies. Scalar particle production in Schwarzschild and Rindler metric. *Journal of Physics*, A8:609–616, 1975. See also [Rindler, 1966]; [Unruh, 1976]; [Davies, 1978]; [Hawking and Page, 1983]; [Unruh and Wald, 1984]; [Wald, 1994]; [Clifton and Halverson, 2001].
- [Davies, 1978] P.C.W. Davies. Thermodynamics of black holes. *Reports on Progress in Physics*, 41:1313–1355, 1978.
- [de Finetti, 1937] B. de Finetti. La prévision, ses lois logiques, et ses sources objectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937. See also [de Finetti, 1972].
- [de Finetti, 1972] B. de Finetti. *Probability, induction and statistics*. Wiley, New York, 1972.
- [Debye, 1912] P. Debye. Zur Theorie der spezifischen Wärme. *Annalen der Physik*, 39:789–839, 1912.
- [Derezinski and Früboes, 2005] J. Derezinski and R. Früboes. Stationary van Hove limit. *Journal of Mathematical Physics*, 46:1–15, 2005.
- [Dickson, 2006] M. Dickson. This volume, chapter 4, 2006.
- [Dixmier, 1957] J. Dixmier. *Les algèbres d'opérateurs dans l'espace Hilbertien (Algèbres de von Neumann)*. Gauthier–Villars, Paris, 1957.
- [Domb and Green, 1972] C. Domb and M.S. Green, editors. *Phase transitions and critical phenomena, I: exact results*. Academic Press, London, New York, 1972.
- [Donkor *et al.*, 2003] E. Donkor, A.R. Pirich and H.E. Brandt, editors. *Quantum information and computation*, International Society for Optical Engineering (SPIE), (2003).
- [Dubin and Sewell, 1970] D.A. Dubin and G.L. Sewell. Time-translations in the algebraic formulation of statistical mechanics. *Journal of Mathematical Physics*, 11:2990–2998, 1970.
- [Dunford and Schwartz, 1964] N. Dunford and J.T. Schwartz. *Linear operators, part i; General theory*. Interscience, New York, 1964.
- [Earman, 2002] J. Earman. What time-reversal invariance and why it matters. *International Journal for the Philosophy of Science*, 16:245–264, 2002.
- [Earman and Rédei, 1996] J. Earman and M. Rédei. Why ergodic theory does not explain the success of equilibrium statistical mechanics. *British Journal for the Philosophy of Science*, 47:63–78, 1996.
- [Eckman *et al.*, 1999] J.P. Eckmann, C.A. Pillet and L. Rey–Bellet. Non-equilibrium statistical mechanics of anharmonic chains coupled to two baths at different temperatures. *Communications in Mathematical Physics*, 201:657–697, 1999.
- [Ehrenfest and Ehrenfest, 1911] P. Ehrenfest and T. Ehrenfest. *Begriffliche Grundlagen der statistischen Auffassungen in der Mechanik*, volume 4/22 of *Enzyklopädie der mathematischen Wissenschaften*, Teubner, Leipzig, 1911.
- [Einstein, 1905a] A. Einstein. Ueber einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt. *Annalen der Physik*, 17:132–148, 1905; transl. in [Stachel, 1998] pp. 177–197. See also [Einstein, 1906a].
- [Einstein, 1905b] A. Einstein. Ueber die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik*, 17:549–560, 1905; transl. in [Einstein, 1926] pp. 1–18; and in [Stachel, 1998] pp. 85–98.
- [Einstein, 1906a] A. Einstein. Zur Theorie der Lichtezeugung und Lichtabsorption. *Annalen der Physik*, 20:190–206, 1906.
- [Einstein, 1906b] A. Einstein. Eine neue Bestimmung der Moleküldimensionen. *Annalen der Physik*, 19:289–306, 1906; transl. in [Einstein, 1926] pp. 36–61; compare with Einstein's Ph.D. dissertation; transl. in [Stachel, 1998] pp. 45–63. For an update of the 1906 paper, see [Einstein, 1911a].

- [Einstein, 1906c] A. Einstein. Zur Theorie der Brownschen Bewegung. *Annalen der Physik*, 19:371-381, 1906; transl. in [Einstein, 1926] pp. 19-35.
- [Einstein, 1907] A. Einstein. Die Plancksche Theorie der Strahlung und die Theorie der spezifischen W rme. *Annalen der Physik*, 22:180-190, 1907.
- [Einstein, 1909a] A. Einstein. Zum gegenw rtigen Stand des Strahlungsproblems. *Physikalische Zeitschrift*, 10:185-193, 1909. See also [Einstein, 1909b].
- [Einstein, 1909b] Ueber die Entwicklung unserer Anschauungen  ber das Wesen und die Konstitution der Strahlung. *Physikalische Zeitschrift*, 10:817-825, 1909.
- [Einstein, 1911a] Berichtigung zu meiner Arbeit. *Annalen der Physik*, 34:591-592, 1911; transl. in [Einstein, 1926] p. 62.
- [Einstein, 1911b] A. Einstein. Elementare Betrachtungen  ber die thermische Molekularbewegung in festen K rpern. *Annalen der Physik*, 35:679-694, 1911.
- [Einstein, 1921] A. Einstein. *Geometry and experience*, (Lecture before the Prussian Academy of Sciences, January 27, 1921) pages 232-249 in [Einstein, 1954].
- [Einstein, 1924] A. Einstein. Quantentheorie des einatomigen idealen Gases. *Sitzungsberichte der Deutschen Akademie der Wissenschaften, Berlin*, 1924:261-267.
- [Einstein, 1926] A. Einstein. *Investigations on the theory of the Brownian movement*, R. F rth, editor, A.D. Cowper, transl., Methuen, 1926, reprinted by Dover, New York, 1956; this is a collection of five papers published by Einstein in the years 1905-1908.
- [Einstein, 1933] A. Einstein. *On the method of theoretical physics*, Herbert Spencer Lecture, 1933; in [Einstein, 1954] pp. 270-276.
- [Einstein, 1954] A. Einstein. *Ideas and Opinions*, S. Bargmann, transl. and rev., Bonanza Books, New York, 1954; based *inter al.* on *Mein Weltbild*, C. Seelig, editor. Querido Verlag, Amsterdam, 1934.
- [Emch, 1964] G.G. Emch. Coarse-graining in Liouville space and master equations. *Helvetica Physica Acta*, 37:532-544, 1964.
- [Emch, 1972a] G.G. Emch. *Algebraic methods in statistical mechanics and quantum field theory*. Wiley, New-York, 1972.
- [Emch, 1972b] G.G. Emch. *The C^* -algebraic approach to phase transitions*, pages 137-175 in [Domb and Green, 1972].
- [Emch, 1976] G.G. Emch. Generalized K-Flows. *Communications in Mathematical Physics*, 49:191-215, 1976. See also [K ummerer and Schr der, 1983].
- [Emch, 1998a] G.G. Emch. *On the need to adapt de Finetti's probabilistic interpretation to QM*, pages 157-166 in [Alicki *et al.*, 1998].
- [Emch, 1998b] G.G. Emch. *Chaotic dynamics in non-commutative geometry*, pages 193-205 in [Strasburger *et al.*, 1998].
- [Emch, 2003] G.G. Emch. *Quantum measurement: a bridge to classical physics*, volume 5105, pages 255-264 in [Donkor *et al.*, 2003].
- [Emch, 2005] G.G. Emch. Probabilistic issues in statistical mechanics. *Studies in History and Philosophy of Modern Physics*, 36:303-322, 2005. See already [Emch, 1998a].
- [Emch and Guenin, 1966] G.G. Emch and M. Guenin. Gauge invariant formulation of the BCS model. *Journal of Mathematical Physics*, 7:915-921, 1966.
- [Emch *et al.*, 1970] G.G. Emch, H.J.F. Knops and E.J. Verboven. Breaking of the Euclidean symmetry with an application to the theory of crystallization. *Journal of Mathematical Physics*, 11:1655-1668, 1970.
- [Emch and Knops, 1970] G.G. Emch and H.J.F. Knops. Pure thermodynamical phases as extremal KMS states. *Journal of Mathematical Physics*, 11:3008-3018, 1970.
- [Emch and Radin, 1971] G.G. Emch and C. Radin. Relaxation of local thermal deviation from equilibrium. *Journal of Mathematical Physics*, 12:2043-2046, 1971.
- [Emch *et al.*, 1994a] G.G. Emch, H. Narnhofer, W. Thirring and G.L. Sewell. Anosov actions on non-commutative algebras. *Journal of Mathematical Physics*, 35:5582-5599, 1994.
- [Emch *et al.*, 1994b] G.G. Emch, G.C. Hegerfeldt and L. Streit, editors. *On Klauder's path: a field trip*. World Scientific, Singapore, 1994.
- [Emch and Liu, 2002] G.G. Emch and C. Liu. *The logic of thermostistical physics*. Springer, Heidelberg, 2002.
- [Erd lyi, 1953] A. Erd lyi. *Higher transcendental functions*. McGrawHill, New York, 1953.
- [Fannes and Tuyls, 1999] M. Fannes and P. Tuyls. A continuity property of quantum dynamical entropy. *Infinite dimensional analysis, quantum probability and related topics*. 2:511-527, 1999.

- [Fermi, 1926] E. Fermi. Zur Quantelung des idealen einatomigen Gases. *Zeitschrift für Physik*, 36:902-912, 1926.
- [Fermi, 1927] E. Fermi. Uno modello statistico per la determinazione d'alcune proprietà dell'atomo. *Atti della Accademia Nazionale dei Lincei, Rendiconti*, 6:602-607, 1927. See also [Fermi, 1928].
- [Fermi, 1928] E. Fermi. Eine statistische Methode zur Bestimmung einiger Eigenschaften des Atomes und ihre Anwendungen auf die Theorie des periodischen Systems der Elemente. *Zeitschrift für Physik*, 48:73-79, 1928.
- [Feynman, 1998] R.P. Feynman. *Statistical mechanics, a set of lectures*. Addison Wesley, Reading MA, 1998.
- [Fokas *et al.*, 2001] A. Fokas, A. Grigoryan, T. Kibble and B. Zegarlinski, editors. *Mathematical Physics 2000*. International Press of Boston, Somerville MA, 2001.
- [Ford *et al.*, 1965] G.W. Ford, M. Kac and P. Mazur. Statistical mechanics of assemblies of coupled oscillators. *Journal of Mathematical Physics*, 6:504-515, 1965.
- [Fowler, 1926] R.H. Fowler. On dense matter. *Monthly Notices Royal Astronomical Society*, 87:114, 1926.
- [Fredenhagen, 2003] K. Fredenhagen. *Locally covariant quantum theory*, in *Proceedings of the International Conference on Mathematical Physics*, Lisbon, 2003.
- [Fredenhagen and Haag, 1990] K. Fredenhagen and R. Haag. On the derivation of Hawking radiation associated with the formation of a black hole. *Communications in Mathematical Physics*, 127:273-284, 1990.
- [Friedrichs, 1953] K.O. Friedrichs. *Mathematical aspects of the quantum theory of fields*. Wiley, New York, 1953.
- [Gamov, 1928] G. Gamow. Zur Quantentheorie des Atomkernes. *Zeitschrift für Physik*, 51:204-212, 1928.
- [Gamov, 1940] G. Gamow. *Mr. Tompkins in Wonderland*. Macmillan, New York, 1940.
- [Garcia-Bondia *et al.*, 2003] J.M. Garcia-Bondia, J. Varilly and H. Figueroa. *Elements of non-commutative geometry*. Birkhäuser, Boston, 2001.
- [Gibbs, 1902] J.W. Gibbs. *Elementary principles of statistical mechanics, developed with special reference to the rational foundation of thermodynamics*. Scribner, New York, 1902.
- [Glimm *et al.*, 1990] J. Glimm, J. Impagliazzo and I. Singer, editors. The legacy of John von Neumann. *Proceedings of Symposia in Pure Mathematics*, volume 50, 1990.
- [Gorini *et al.*, 1978] V. Gorini, A. Frigerio, M. Verri, A. Kossakowski and E.C.G. Sudarshan. Properties of quantum Markovian master equations. *Reports on Mathematical Physics*, 13:149-173, 1978.
- [Gouy, 1888] L. Gouy. Note sur le mouvement brownien. *Journal de Physique (2)*, 7:561-564, 1888.
- [Grad, 1958] H. Grad. *Principles of the kinematic theory of gases*, volume 12 of *Handbuch der Physik*. Springer, Berlin, 1958. See also [Grad, 1963].
- [Grad, 1963] H. Grad. Asymptotic theory of the Boltzmann equation. *Physics of Fluids*, 6:147-181, 1963.
- [Grad, 1967] H. Grad. *Levels of description in statistical mechanics and thermodynamics*, pages 49-76 in [Bunge, 1967].
- [Greenleaf, 1969] F.P. Greenleaf. *Invariant mean on topological groups*. van Nostrand-Reinhold, New York, 1969.
- [Gurney and Condon, 1928] R.W. Gurney and E.U. Condon. Quantum mechanics and radioactive disintegration. *Nature*, 122:439, 1928. See also [Gurney and Condon, 1929].
- [Gurney and Condon, 1929] R.W. Gurney and E.U. Condon. Quantum mechanics and radioactive disintegration. *Physical Review*, 33:127-140, 129. See also [Gamov, 1928].
- [Gutzwiller, 1990] M.C. Gutzwiller. *Chaos in classical and quantum mechanics*. Springer, New York, 1990.
- [Haag, 1955] R. Haag. On quantum field theory. *Matematisk-Fysiske Meddelelser Danske Videnskabernes Selskab*, 29(12), 1955.
- [Haag, 1959a] R. Haag. *Discussion des "axiomes" et des propriétés asymptotiques d'une théorie des champs locale avec particules composées*, in *Les problèmes mathématiques de la théorie quantique des champs*, Colloque CNRS, Lille, 1958; CNRS, Paris, 1959.
- [Haag, 1959b] R. Haag. The framework of quantum field theory. *Il Nuovo Cimento*, 14:131-152, 1959.

- [Haag, 1962] R. Haag. The mathematical structure of the Bardeen–Cooper–Schrieffer model. *II Nuovo Cimento*, 25:287–299, 1962.
- [Haag and Schroer, 1962] R. Haag and B. Schroer. Postulates of quantum field theory. *Journal of Mathematical Physics*, 3:248–256, 1962.
- [Haag and Kastler, 1964] R. Haag and D. Kastler. An algebraic approach to quantum field theory. *Journal of Mathematical Physics*, 5:848–861, 1964.
- [Haag *et al.*, 1967] R. Haag, N. Hugenholtz and M. Winnink. On the equilibrium states in quantum statistical mechanics. *Communications in Mathematical Physics*, 5:215–236, 1967.
- [Haag *et al.*, 1974] R. Haag, D. Kastler and E.B. Trych–Pohlmeyer. Stability and equilibrium states. *Communications in Mathematical Physics*, 38:173–193, 1974. See [Haag and Trych–Pohlmeyer, 1977].
- [Haag and Trych–Pohlmeyer, 1977] R. Haag and E.B. Trych–Pohlmeyer. Stability properties of equilibrium states. *Communications in Mathematical Physics*, 56:213–224, 1977.
- [Haag *et al.*, 1994] R. Haag, H. Narnhofer and U. Stein. On quantum field theory in gravitational background. *Communications in Mathematical Physics*, 94:219–238, 1994.
- [Halvorson, 2006] H. Halvorson. This volume, chapter 8, 2006.
- [Hawking, 1975] S.W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43:199–220, 1975.
- [Hawking, 1981] R.W. Hawking. Interacting quantum fields around a black hole. *Communications in Mathematical Physics*, 80:421–442, 1981.
- [Hawking and Page, 1983] S.W. Hawking and D.N. Page. Thermodynamics of black holes in anti-de Sitter space. *Communications in Mathematical Physics*, 87:577–588, 1983.
- [Heisenberg, 1928] W. Heisenberg. Zur Theorie der Ferromagnetismus. *Zeitschrift f r Physik*, 49:419–436, 1928.
- [Hepp, 1972] K. Hepp. Quantum theory of measurement and macroscopic observables. *Helvetica Physica Acta*, 45:237–248, 1972.
- [Hida, 1970] T. Hida. *Stationary stochastic processes*. Princeton University Press, Princeton NJ, 1970.
- [Hilbert, 1899] D. Hilbert. *Grundlagen der Geometrie*. Teubner, Leipzig, 1899; translated by L. Unger: *Foundations of Geometry*. Open Court, La Salle IL, 1971.
- [Hilbert, 1900] D. Hilbert. *Sur les probl mes futurs des math matiques*, pages 58–114 in *Comptes rendus du deuxi me congr s international des math maticiens, Paris, 1900*, Gauthier–Villars, Paris, 1902. See also [Browder, 1976].
- [Hilbert, 1918] D. Hilbert. Axiomatisches Denken. *Mathematische Annalen*, 78:405–415, 1918.
- [Ho *et al.*, 1993] T.G. Ho, L.J. Landau and A.J. Wilkins. On the weak coupling limit for a Fermi gas in a random potential. *Reviews in Mathematical Physics*, 5:209–298, 1993.
- [Huang, 1965] K. Huang. *Statistical mechanics*. Wiley, New York, 1965.
- [Hudson, 1998] R.L. Hudson. *Quantum stochastic analysis after four decades*, pages 194–217 in *Probability towards 2000*, Springer Lecture Notes in Statistics 128, New York, 1998.
- [Ingarden *et al.*, 1999] R.S. Ingarden, R. Smolukowski, S. Chandrasekhar and M. Kac. *Marian Smolukowski: his life and scientific work*. Polish Scientific Publishers PWN, Warsaw, 1999.
- [It , 1987] K. It , editor. *Encyclopedic dictionary of mathematics*. The MIT Press, Cambridge MA and London England, 1987; 2nd edition of the English translation of the 3rd Japanese edition.
- [Jammer, 1966] M. Jammer. *The conceptual development of quantum mechanics*. McGraw–Hill, New York, 1966.
- [Jammer, 1974] M. Jammer. *The philosophy of quantum mechanics*. Wiley, New York, 1974.
- [Jaynes, 1967] E.T. Jaynes. *Foundations of probability theory and statistical mechanics*, pages 77–101 in [Bunge, 1967]. See also [Jaynes, 1989].
- [Jaynes, 1989] *Papers on probability, statistics and statistical mechanics*. Kluwer, Dordrecht, 1989. 1st ed. 1983.
- [Jeans, 1905a] J.B. Jeans. On the partition of energy between matter and ether. *Philosophical Magazine*, 10:91–98, 1905. See also [Jeans, 1905b].
- [Jeans, 1905b] J.B. Jeans. Comparison between two theories of radiation. *Nature*, 72:293, 1905.
- [Josephson, 1982] B.D. Josephson. *The discovery of tunnelling supercurrents*, pages 157–164 in [Lundqvist, 1982].
- [Kac, 1968] M. Kac. *Mathematical mechanisms of phase transitions*, pages in 241–305 in volume i of [Chr tien *et al.*, 1968].

- [Kadison, 1990] R.V. Kadison. *Operator algebras — an overview*, pages 61–89 in [Glimm *et al.*, 1990].
- [Kadison and Ringrose, 1983/1986] R.V. Kadison and J.R. Ringrose. *Fundamentals in the theory of operator algebras, i. Elementary theory; ii. Advanced theory*. Wiley, New York, 1983/1986.
- [Kastler, 1976] D. Kastler. Equilibrium states of matter and operator algebras. *Istituto Nazionale di Alta Matematica, Symposia Mathematica*, 20:49–107, 1976. See also [Haag *et al.*, 1974] and [Narnhofer and Thirring, 1982].
- [Kastler and Robinson, 1966] D. Kastler and D.W. Robinson. Invariant states in statistical mechanics. *Communications in Mathematical Physics*, 3:151–180, 1966. See also [Robinson and Ruelle, 1967].
- [Kasler *et al.*, 1967] D. Kastler, R. Haag and L. Michel. Central decomposition of ergodic states. *Séminaire de physique théorique, Marseille*, 1967. See also [Kastler *et al.*, 1972], [Kastler, 1976], [Kastler and Robinson, 1966], [Robinson and Ruelle, 1967].
- [Kastler *et al.*, 1972] D. Kastler, M. Mebkhout, G. Loupias and L. Michel. Central decomposition of invariant states with applications to the group of time translations and of Euclidean transformations in algebraic field theory. *Communications in Mathematical Physics*, 27:195–222, 1972.
- [Kay, 1993] B.S. Kay. Sufficient conditions for quasifree states on spacetimes with a bifurcate Killing horizon. *Journal of Mathematical Physics*, 34:4519–4539, 1993.
- [Kay and Wald, 1991] B.S. Kay and R.M. Wald. Theorems on the uniqueness and thermal properties of stationary nonsingular, quasifree states on spacetimes with a bifurcate Killing horizon. *Physics Reports*, 207:49–136, 1991. See also [Kay, 1993].
- [Khinchin, 1957] A. Khinchin. *Mathematical foundations of information theory*. Dover, New York, 1957.
- [Khrenninkov, 2005] A. Khrenninkov, editor. *Foundations of probability and physics*. American Institute of Physics, New York, 2005.
- [Kishimoto, 1976] A. Kishimoto. On uniqueness of KMS states of one-dimensional quantum lattice systems. *Communications in Mathematical Physics*, 47:167–170, 1976.
- [Kolmogorov, 1933] A.N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933; translated by N. Morrison: *Foundations of the theory of probability*. Chelsea, New York, 1956.
- [Kossakowski *et al.*, 1977] A. Kossakowski, A. Frigerio, V. Gorini and M. Verri. Quantum detailed balance and the KMS condition. *Communications in Mathematical Physics*, 57:97–110, 1977; Erratum [Kossakowski *et al.*, 1978].
- [Kossakowski *et al.*, 1978] A. Kossakowski, A. Frigerio, V. Gorini and M. Verri. Erratum: “Quantum detailed balance and the KMS condition”. *Communications in Mathematical Physics*, 60:96, 1978.
- [Kramers and Wannier, 1941] H.A. Kramers and G.H. Wannier. Statistics of the two-dimensional ferromagnet, i, ii. *Physical Review*, 60:252–262, 263–276, 1941.
- [Kubo, 1957] R. Kubo. Statistical mechanics of irreversible processes. *Journal of the Physical Society of Japan*, 12:570–586, 1957.
- [Kümmerer and Schröder, 1983] B. Kümmerer and W. Schröder. A Markov dilation of a non-quasifree Bloch evolution. *Communications in Mathematical Physics*, 90:251–262, 1983.
- [Landau, 1932] L.D. Landau. On the theory of stars. *Physikalische Zeitschrift der Sowjetunion*, 1:285–288, 1932.
- [Landau and Lifshitz, 1958a] L.D. Landau and E.M. Lifshitz, *Quantum mechanics*. Pergamon, London, 1958.
- [Landau and Lifshitz, 1958b] L.D. Landau and E.M. Lifshitz, *Statistical physics*. Pergamon, London 1958.
- [Landsman, 2006] N. Landsman. This volume, Chapter 5, 2006.
- [Le Bris and Lions, 2005] C. Le Bris and P.–L. Lions. From atoms to crystals: a mathematical journey. *Bulletin of the American Mathematical Society*, 42:291–363, 2005.
- [Lenard, 1970] A. Lenard, editor. *Statistical mechanics and mathematical problems*, volume 20 of *Lecture Notes in Physics*. Springer, New York, 1970.
- [Lieb, 1982a] E.H. Lieb. Thomas-Fermi and related theories of atoms and molecules. *Reviews of Modern Physics*, 53:603–641, 1982; reprinted in [Lieb, 1997] pp. 259–297. See also [Lieb, 1982b].

- [Lieb, 1982b] E.H. Lieb. Erratum: “Thomas-Fermi and related theories of atoms and molecules”. *Reviews of Modern Physics*, 54:311, 1982; reprinted in [Lieb, 1997] p. 298.
- [Lieb, 1990] E.H. Lieb. The stability of matter: from atoms to stars. *Bulletin of the American Physical Society*, 22:1-49, 1990; reprinted in [Lieb, 1997] pp. 11-59.
- [Lieb, 1997] E.H. Lieb. *Selecta: The stability of matter, from atoms to stars*. Springer, Berlin, Heidelberg, 1997.
- [Lieb, 2001] E.H. Lieb. *The Bose-gas: a subtle many-body problem*, pages 91-111 in [Fokas et al., 2001].
- [Lieb and Simon, 1977] E.H. Lieb and B. Simon. Thomas-Fermi theory of atoms, molecules and solids. *Advances in Mathematics*, 23:22-116, 1977.
- [Liu and Emch, 2005] C. Liu and G.G. Emch. Explaining quantum spontaneous symmetry breaking. *Studies in History and Philosophy of Modern Physics*, 36:137-163, 2005.
- [Lowe and Nordberg, 1957] I.J. Lowe and R.E. Nordberg. Free-induction decay in solids. *Physical Review*, 107:46-61, 1957.
- [Lummer and Pringsheim, 1900] O. Lummer and E. Pringsheim. Ueber die Strahlung der schwarzen K rper f r lange Wellen. *Verhandlungen der Deutschen Physikalischen Gesellschaft*, 2:163-180, 1900.
- [Lundqvist, 1982] S. Lundqvist, editor. *Nobel lectures 1971-1980*. World Scientific, Singapore 1982.
- [Martin and Schwinger, 1959] P.C. Martin and J. Schwinger. Theory of many-particle systems. *Physical Review*, 115:1342-1373, 1959.
- [Martin, 1979] Ph.A. Martin. *Mod les en m canique des processus irr versibles*, volume 103 of *Lecture Notes in Physics*. Springer, Berlin, 1979. See also [Martin and Emch, 1975] and [Spohn, 1977].
- [Martin and Emch, 1975] Ph.A. Martin and G.G. Emch. A rigorous model sustaining van Hove’s phenomenon. *Helvetica Physica Acta*, 48:48-59, 1975. See also [Ho et al., 1993].
- [Messiah, 1960] A. Messiah. *M canique quantique*. Dunod, Paris, 1960.
- [Michel and Ruelle, 1970] L. Michel and D. Ruelle, editors. *Syst mes   un nombre infini de degr s de libert *, CNRS, Paris, 1970.
- [Murray, 1990] F.J. Murray. *The rings of operators papers*, pages 57-60 in [Glimm et al., 1990].
- [Murray and von Neumann, 1936] F.J. Murray and J. von Neumann. Rings of Operators. *Annals of Mathematics*, 37:116-229, 1936.
- [Murray and von Neumann, 1937] F.J. Murray and J. von Neumann. Rings of Operators ii. *Transactions of the American Mathematical Society*, 41:208-248, 1937.
- [Narnhofer, 1998] H. Narnhofer. *Control of weak asymptotic abelianness with the help of the cross-product construction*, pages 331-339 in [Alicki et al., 1998].
- [Narnhofer, 2001] H. Narnhofer. Kolmogorov systems and Anosov systems in quantum theory. *Infinite Dimensional Analysis, Quantum Probability and Related Topics*, 4:85-119, 2001.
- [Narnhofer, 2005] H. Narnhofer. Roots for quantum ergodic theory in von Neumann’s work. *Reports on Mathematical Physics*, 55:93-107, 2005.
- [Narnhofer and Thirring, 1982] H. Narnhofer and W. Thirring. Adiabatic theorem in quantum statistical mechanics. *Physical Review A*, 26:3646-3652, 1982. See also [Thirring, 1983a].
- [Narnhofer and Thirring, 1989] H. Narnhofer and W. Thirring. Quantum K-systems. *Communications in Mathematical Physics*, 125:565-577, 1989.
- [Narnhofer and Thirring, 1994a] H. Narnhofer and W. Thirring. Clustering for algebraic K-systems. *Letters in Mathematical Physics*, 30:307-316, 1994.
- [Narnhofer and Thirring, 1994b] H. Narnhofer and W. Thirring. *Dynamical entropy of quantum systems and their Abelian counterpart*, pages 127-145 in [Emch et al., 1994b].
- [Ojima, 2003] I. Ojima. *How to formulate non-equilibrium local states in QFT - general characterization and extension to curved space-time*, pages 365-384 in [Arafune et al., 2003].
- [Parthasarathy, 1995] K.R. Parthasarathy. *Quantum stochastic calculus*, pages 1024-1035 in *Proceedings, International Congress of Mathematicians, Zurich, 1994*, Birkh user, Basel, 1995.
- [Paschen and Wanner, 1899] F. Paschen and H. Wanner. Eine photometrische Methode zur Bestimmung der Exponentialkonstanten der Emissionsfunktion. *Berliner Berichte*, 1899:5-11.
- [Peierls, 1936] R. Peierls. On Ising’s model of ferromagnetism. *Proceedings of the Cambridge Philosophical Society*, 32:477-481, 1932.
- [Peierls, 1979] R. Peierls. *Surprises in theoretical physics*. Princeton University Press, Princeton NJ, 1979.

- [Pitaevskii and Stringari, 2003] L. Pitaevskii and S. Stringari. *Bose–Einstein condensation*. Oxford University Press, Oxford, 2003.
- [Planck, 1900a] M. Planck. Ueber eine Verbesserung der Wienschen Spektralgleichung. *Verhandlungen der Deutschen Physikalischen Gesellschaft*, 2:202–204, 1900.
- [Planck, 1900b] M. Planck. Zur Theorie des Gesetzes der Energieverteilung im Normalspektrum. *Verhandlungen der Deutschen Physikalischen Gesellschaft*, 2:237–245, 1900; *ibid. Annalen der Physik*, 4:553–563, 1901.
- [Planck, 1911] M. Planck. Deux récentes théories thermodynamiques: théorème de Nernst et hypothèse des quanta (Conférence prononcée le 16 décembre 1911 à la Société chimique allemande de Berlin); appended to [Planck, 1913].
- [Planck, 1913] M. Planck. *Leçons de thermodynamique*, 3rd ed., Hermann, Paris 1913.
- [Prigogine, 1997] I. Prigogine. *The end of certainty: time, chaos and the new laws of nature*. The Free Press, New York, 1997.
- [Pusz and Woronowicz, 1978] W. Pusz and S.L. Woronowicz. Passive states and KMS states for general quantum systems. *Communications in Mathematical Physics*, 58:273–290, 1978.
- [Rauch, 2005] H. Rauch. *Quantum phenomena tested by neutron interferometry*, pages 331–339 in [Khrenninkov, 2005].
- [Rauch and Werner, 2000] H. Rauch and S.A. Werner. *Neutron interferometry*, Oxford Scientific Publishers, Oxford, 2000.
- [Rayleigh, 1900] J.W.S. Rayleigh. Remarks upon the law of complete radiation. *Philosophical Magazine*, 49:539–540, 1900.
- [Rayleigh, 1902] J.W.S. Rayleigh. The constant of radiation as calculated from molecular data. *Nature*, 72:243, 1902.
- [Rindler, 1966] W. Rindler. Kruskal space and the uniformly accelerated frame. *American Journal of Physics*, 34:1174–1178, 1966.
- [Robinson, 1965] D.W. Robinson. A theorem concerning positive metric. *Communications in Mathematical Physics*, 1:89–94, 1965.
- [Robinson, 1971] D.W. Robinson. *The thermodynamical pressure in quantum statistical mechanics*, volume 9 of *Lecture Notes in Physics*. Springer, Berlin, 1971.
- [Robinson, 1973] D.W. Robinson. Return to equilibrium. *Communications in Mathematical Physics*, 31:171–189, 1973.
- [Robinson and Ruelle, 1967] D.W. Robinson and D. Ruelle. Extremal invariant states. *Annales de l'Institut Henri Poincaré*, 6:299–310, 1967. See already [Kastler and Robinson, 1966].
- [Ruelle, 1968a] D. Ruelle. *Statistical mechanics*. Benjamin, New York, 1968.
- [Ruelle, 1968b] D. Ruelle. Statistical mechanics of a one-dimensional lattice gas. *Communications in Mathematical Physics*, 9:267–278, 1968.
- [Sakai, 1971] S. Sakai. *C^* -algebras and W^* -algebras*. Springer, New York, 1971.
- [Schiff, 1955] L.I. Schiff. *Quantum mechanics*, 2nd ed., MacGraw Hill, London, 1955.
- [Schrieffer, 1974] J.R. Schrieffer. *Theory of superconductivity*. Benjamin, New York, 1974.
- [Schroer and Wiesbrock, 2000] B. Schroer and H.W. Wiesbrock. Modular theory and geometry. *Reviews in Mathematical Physics*, 12:301–326, 2000.
- [Segal, 1990] A.F. Segal. *Paul the convert: the apostolate and apostasy of Saul the Pharisee*. Yale University Press, New Haven, 1990.
- [Segal, 1947] I.E. Segal. Postulates for general quantum mechanics. *Annals of Mathematics*, 48:930–948, 1947.
- [Segal, 1958] I.E. Segal. Distributions in Hilbert space and canonical systems of operators. *Transactions of the American Mathematical Society*, 88:12–41, 1958.
- [Segal, 1959] I.E. Segal. Foundations of the theory of dynamical systems of infinitely many degrees of freedom, i. *Matematisk-Fysiske Meddelelser Danske Videnskabernes Selskab*, 31:2, 1959.
- [Segal, 1961] I.E. Segal. Foundations of the theory of dynamical systems of infinitely many degrees of freedom, ii. *Canadian Journal of Mathematics*, 13:1–18, 1961.
- [Segal, 1962] I.E. Segal. Foundations of the theory of dynamical systems of infinitely many degrees of freedom, iii. *Journal of Mathematics*, 6:500–523, 1962.
- [Segal, 1963] I.E. Segal. *Mathematical problems of relativistic physics*. Lecture Notes in Applied Mathematics. American Mathematical Society, Providence RI, 1963.
- [Sewell, 1973] G.L. Sewell. States and dynamics of infinitely extended physical systems. *Communications in Mathematical Physics*, 33:42–51, 1973.

- [Sewell, 1977] G.L. Sewell. KMS conditions and local thermodynamical stability of quantum lattice systems, ii. *Communications in Mathematical Physics*, 55:53–61, 1977.
- [Sewell, 1980a] G.L. Sewell. Relativity of temperature and the Hawking effect. *Physics Letters A*, 79:23–24, 1980. See also [Sewell, 1982a].
- [Sewell, 1980b] G.L. Sewell. Stability, equilibrium and metastability in statistical mechanics. *Physics Reports*, 57:307–342, 1980.
- [Sewell, 1982a] G.L. Sewell. Quantum fields on manifolds: PCT and gravitationally induced thermal states. *Annals of Physics*, 141:201–224, 1982.
- [Sewell, 1982b] G.L. Sewell. W^* -dynamics of infinite quantum systems. *Letters in Mathematical Physics*, 6:209–213, 1982. See already [Dubin and Sewell, 1970] and [Sewell, 1973].
- [Sewell, 2002] G.L. Sewell. *Quantum mechanics and its emergent macrophysics*. Princeton University Press, Princeton and Oxford, 2002.
- [Sewell, 2005] G.L. Sewell. On the mathematical structure of quantum measurement theory. *Reports on Mathematical Physics*, 56:271–290, 2005.
- [Shapiro and Teukolsky, 1983] S.L. Shapiro and S.A. Teukolsky. *Black holes, white dwarfs, and neutron stars: the physics of compact objects*. Wiley, New York, 1983.
- [Smolukowski, 1906a] M. Smolukowski. Sur le chemin moyen parcouru par les mol cules d’un gas et sur son rapport avec la th orie de la diffusion. *Bulletin International de l’Acad mie des Sciences de Cracovie*, 1906:202–213; reprinted with translation, pages 29–42 in [Ingarden *et al.*, 1999]. See also [Smolukowski, 1906b].
- [Smolukowski, 1906b] M. Smolukowski. Zur kinetischen Theorie der Brownschen Molekular Bewegung und der Suspensionen. *Annalen der Physik*, 21:755–780, 1906.
- [Smolukowski, 1916] M. Smolukowski. Drei Vortr ge  ber Diffusion, Brownsche Molekularbewegung und Koagulation von Kolloidteilchen. *Physikalische Zeitschrift*, 17:557–571, 585–599, 1916; reprinted with translation, pages 43–127 in [Ingarden *et al.*, 1999].
- [Spohn, 1977] H. Spohn. Derivation of the transport equations for electrons moving through random impurities. *Journal of Statistical Physics*, 17:385–412, 1977.
- [Stachel, 1998] J. Stachel, editor. *Einstein’s miraculous year*. Princeton University Press, Princeton NJ, 1998.
- [Stefan, 1879] J. Stefan. Ueber die Beziehung zwischen der Wärmestrahlung und der Temperatur. *Wiener Berichte*, 79:391–428, 1879.
- [St rmer, 1969] E. St rmer. Symmetric states on infinite tensor products of C^* -algebras. *Journal of Functional Analysis*, 3:48–68, 1969.
- [Strasburger *et al.*, 1998] A. Strasburger, S.T. Ali, J.-P. Antoine, J.-P. Gazeau and A. Odziejewicz, editors. *Quantization, coherent states and Poisson structures. Proceedings of the 1995 Bialowieza Workshop*. Polish Scientific Publications PWN, Warszawa, 1998.
- [Streater, 1967] R.F. Streater. The Heisenberg ferromagnet as a quantum field theory. *Communications in Mathematical Physics*, 6:233–247, 1967.
- [Streater and Wightman, 1964] R.F. Streater and A.S. Wightman. *PCT, spin & statistics, and all that*. Benjamin, New York, 1964.
- [Summers and Verch, 1996] S.J. Summers and R. Verch. Modular inclusion, the Hawking temperature and quantum field theory in curved space–time. *Letters in Mathematical Physics*, 37:145–158, 1996.
- [Summers and White, 2003] S.J. Summers and R.K. White. On deriving space–time from quantum observables and states. *Communications in Mathematical Physics*, 237:203–220, 2003.
- [Szasz, 1996] D. Szasz. Boltzmann’s ergodic hypothesis, a conjecture for centuries. *Studia scientiarum mathematicarum Hungarica*, 31:299–322, 1996.
- [Takeda, 1955] Z. Takeda. Inductive limit and infinite direct product of operator algebras. *Tohoku Mathematical Journal*, 7:68–86, 1955.
- [Takesaki, 1970a] M. Takesaki. *Tomita’s theory of modular algebras and its applications*, volume 128 of *Lecture Notes in Mathematics*. Springer, New York, 1970.
- [Takesaki, 1970b] M. Takesaki. *States and automorphisms of operator algebras – standard representations and KMS boundary condition*, pages 205–246 in [Lenard, 1970].
- [Takesaki, 1970c] M. Takesaki. Disjointness of the KMS states of different temperatures. *Communications in Mathematical Physics*, 17:33–41, 1970.
- [Talman, 1968] J.C. Talman. *Special functions: a group theoretic approach; based on lectures by Eugene P. Wigner*. Benjamin, New York, 1968.
- [Taylor, 1909] G.I. Taylor. Interference fringes in feeble light. *Proceedings of the Cambridge Philosophical Society, Mathematical and Physical Sciences*, 15:114–115, 1909.

- [Thirring, 1968] W. Thirring. On the mathematical structure of the BCS-Model, ii. *Communications in Mathematical Physics*, 7:181–189, 1968. See already [Thirring and Wehrl, 1967].
- [Thirring, 1983a] W. Thirring. On the equivalence of adiabatic invariance and KMS condition. *Foundations of Physics*, 13:409–415, 1983.
- [Thirring, 1983b] W. Thirring. *Quantum mechanics of large systems*. Springer, New York, Wien, 1983.
- [Thirring and Wehrl, 1967] W. Thirring and A. Wehrl. On the mathematical structure of the BCS-Model. *Communications in Mathematical Physics*, 4:303–314, 1967. See also [Thirring, 1968].
- [Thomas, 1927] L.H. Thomas. The calculation of atomic fields. *Proceedings of the Cambridge Philosophical Society*, 23:542–548, 1927.
- [Thorne, 1994] K.S. Thorne. *Black holes & time warps, Einstein's outrageous legacy*, Norton, New York, 1994.
- [Tomita, 1967] M. Tomita. *Standard forms of von Neumann algebras*. 5th Functional Analysis Symposium, Mathematical Society of Japan, Sendai, 1967.
- [Truesdell, 1945] C. Truesdell. A function which occurs in the theory of the structure of polymers. *Annals of Mathematics*, 46:144–157, 1945.
- [Tuyls, 1998] P. Tuyls. *Comparing quantum dynamical entropies*, pages 411–420 in [Alicki et al., 1998]. See already [Alicki and Fannes, 1994], and also [Fannes and Tuyls, 1999].
- [Uffink, 1995] J. Uffink. Can the minimum entropy principle be regarded as a consistency requirement? *Studies in History and Philosophy of Modern Physics*, 26:223–261, 1995.
- [Uffink, 2006] J. Uffink. This volume, Chapter 9, 2006.
- [Uhlenbeck, 1967] G.E. Uhlenbeck. *Discussion comment*, page 56 in [Bak, 1967].
- [Uhlenbeck, 1968] G.E. Uhlenbeck. *An outline of statistical mechanics*, pages 1–29 in volume ii of [Cohen, 1962/1968].
- [Unruh, 1976] W.G. Unruh. Notes on black-hole evaporation. *Physical Review D*, 14:870–892, 1976.
- [Unruh and Wald, 1984] W.G. Unruh and R.M. Wald. What happens when an accelerating observer detects a Rindler particle? *Physical Review D*, 29:1047–1056, 1984.
- [van Hove, 1952] L. van Hove. Les difficultés de divergences sur un modèle particulier de champ quantifié. *Physica*, 18:145–159, 1952.
- [van Hove, 1955] L. van Hove. Quantum mechanical perturbations giving rise to a statistical transport equation. *Physica*, 21:517–540, 1955.
- [van Hove, 1957] L. van Hove. The approach to equilibrium in quantum statistics. *Physica*, 23:441–480, 1957.
- [van Hove, 1959] L. van Hove. The ergodic behaviour of quantum many-body systems. *Physica*, 25:268–276, 1959.
- [van Hove, 1961] L. van Hove. *The problem of master equations*, pages 155–169 in [Caldirola, 1961].
- [van Hove, 1962] L. van Hove. *Master equation and the approach to equilibrium for quantum systems*, pages 157–172 in volume i of [Cohen, 1962/1968].
- [van Kampen, 1954] N.G. van Kampen. Quantum statistics of irreversible processes. *Physica*, 20:603–622, 1954. See also [van Kampen, 1988].
- [van Kampen, 1962] N.G. van Kampen. *Fundamental problems in statistical mechanics of irreversible processes*, pages 173–202 in volume i of [Cohen, 1962/1968].
- [van Kampen, 1967] N.G. van Kampen. *Discussion remark*, page 67 in [Bak, 1967].
- [van Kampen, 1988] N.G. van Kampen. Ten theorems about quantum mechanical measurements. *Physica A*, 153:97–113, 1988.
- [Verboven, 1967] E.J. Verboven. *Quantum thermodynamics of an infinite system of harmonic oscillators*, pages 49–54 in [Bak, 1967].
- [Vilenkin, 1968] N. Vilenkin. *Special functions and the theory of group representations*, volume 22 of American Mathematical Society Translations, Providence RI, 1968.
- [von Mises, 1928] R. von Mises. *Wahrscheinlichkeit, Statistik und ihre Wahrheit*, Springer, Wien, 1928. See also [von Mises, 1957].
- [von Mises, 1957] R. von Mises. *Probability, statistics and truth*, McMillan, New York, 1957; prepared by H. Geringer from the 3rd German ed. (1951).
- [von Neumann, 1932a] J. von Neumann. Proof of the quasi-ergodic theorem. *Proceedings of the National Academy of Sciences (USA)*, 18:70–83, 1932. See also [von Neumann, 1932b].

- [von Neumann, 1932b] J. von Neumann. Physical applications of the ergodic hypothesis. *Proceedings of the National Academy of Sciences (USA)*, 18:263–266, 1932. See also [Birkhoff, 1931].
- [von Neumann, 1932c] J. von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Springer, Berlin, 1932; translated by R.T. Beyer: *Mathematical foundations of quantum mechanics*, Princeton University Press, Princeton NJ, 1955; page references are to this edition.
- [Wald, 1977] R.M. Wald. *Space, time, and gravity: the theory of big bang and black holes*. University of Chicago Press, Chicago, 1977.
- [Wald, 1994] R.M. Wald. *Quantum field theory in curved spacetime and black hole thermodynamics*. University of Chicago Press, Chicago, 1994. For an elementary introduction, see [Wald, 1977].
- [Wannier, 1966] G.H. Wannier. *Statistical physics*. John Wiley, New York, 1966.
- [Weinberg, 1972] S. Weinberg. *Gravitation and cosmology, principles and applications of the general theory of relativity*. Wiley, New York, 1972.
- [Weiss, 1907] P. Weiss. L’hypoth ese du champ mol culaire et la propri et  ferromagn tique. *Journal de Physique*, 6:661–690, 1907.
- [Whittaker and Watson, 1927] A.N. Whittaker and G.N. Watson, *A course in modern analysis*, 4th ed., Cambridge University Press, London, 1927.
- [Whitten-Wolfe and Emch, 1976] B. Whitten-Wolfe and G.G. Emch. A mechanical quantum measuring process. *Helvetica Physica Acta*, 49:45–55, 1976.
- [Wick et al., 1952] G.C. Wick, A.S. Wightman and E.P. Wigner. Intrinsic parity of elementary particles. *Physical Review*, 88:101–105, 1952.
- [Wien, 1894] W. Wien. Temperatur und Entropie der Strahlung. *Annalen der Physik*, 52:132–165, 1894.
- [Wien, 1896] W. Wien. Ueber die Energieverteilung im Emissionspektrum eines schwarzen K rpers. *Annalen der Physik*, 58:662–669, 1896.
- [Wiesbrock, 1997] H.W. Wiesbrock. Half-sided modular inclusions of von Neumann algebras. *Communications in Mathematical Physics*, 184:583–685, 1997. See also [Schroer and Wiesbrock, 2000].
- [Wightman and Schweber, 1955] A.S. Wightman and S. Schweber. Configuration space methods in relativistic quantum field theory. *Physical Review*, 98:812–837, 1955.
- [Wigner, 1931] E.P. Wigner. *Gruppentheorie und ihre Anwendung auf die Quantenmechanik der Atomspektren*. Vieweg, Braunschweig, 1931; translated by J.C. Griffin: *Group theory and its application to the quantum mechanics of atomic spectra*. Academic Press, New York, 1959.
- [Wigner, 1962] E.P. Wigner. *Oral admonition to young researchers*. Private communication (1962), which I understand many others have had to face, before and since.
- [Wigner, 1984] E.P. Wigner. *Review of the quantum-mechanical measurement problem*, pages 63–82 in *Science, computers and the information onslaught*, D.M. Kerr et al., editors, Academic Press, New York, 1984; reprinted in [Wigner, 1995] pp. 225–244.
- [Wigner, 1995] E.P. Wigner. *Philosophical reflections and syntheses*, volume vi of *The collected works of Eugene Paul Wigner*. Springer, Heidelberg, 1995 (soft-cover, 1997). Annotated by G.G. Emch.
- [Wigner, 1997] E.P. Wigner. *Foundations of quantum mechanics*, part ii of volume iii of *The collected works of Eugene Paul Wigner*. Springer, Heidelberg, 1997. Annotated by A. Shimony.
- [Winnink, 1967] M. Winnink. *On the equilibrium states of infinite quantum systems at $T \neq 0$* , pages 59–66 in [Bak, 1967].
- [Zeh, 1970] H.D. Zeh. On the interpretation of measurement in quantum theory. *Foundations of Physics*, 1:69–76, 1970.

ISSUES IN THE PHILOSOPHY OF COSMOLOGY

George F R Ellis

1 INTRODUCTION

Cosmology is the study of the large-scale structure of the Universe, where ‘*the Universe*’ means all that exists in a physical sense [Harrison, 2000]. This is to be distinguished from *the Observable Universe*, namely that part of the Universe containing matter accessible to our astronomical observations, which is a subset of the Universe proper. Thus cosmology considers the vast domain of galaxies, clusters of galaxies, quasi-stellar objects, etc., observable in the sky by use of telescopes of all kinds, examining their nature, distribution, origins, and relation to their larger environment. *Observational cosmology* [Hoyle, 1960; Kristian and Sachs, 1966; Gunn *et al.*, 1978; Sandage *et al.*, 1993; Bothun, 1998] aims to determine the large-scale geometry of the observable universe and the distribution of matter in it from observations of radiation emitted by distant objects, while *physical cosmology* [Peebles, 1971; Sciama, 1971; Weinberg, 1972; Silk, 2001; Perkins, 2005; Dodelson, 2003] is the study of interactions during the expansion of the universe in its early hot big bang phase, and *astrophysical cosmology* [Sciama, 1971; Peebles, 1993b; Padmanbhan, 1993; Rees, 1995; Dodelson, 2003] studies the resulting later development of large-scale structures such as galaxies and clusters of galaxies. Various forms of quantum cosmology (see e.g. [Hawking, 1993; Gibbons *et al.*, 2003; Copeland *et al.*, 2005]) and studies of particle physics aspects of cosmology [Kolb and Turner, 1990; Peacock, 1999; Allday, 2002; Perkins, 2005; Dodelson, 2003] attempt to characterize the epochs before the hot big bang phase. These studies function in a mainly symbiotic way, each informing and supplementing the others to create an overall cosmological theory of the origin and evolution of the physical universe [Bondi, 1960; Harrison, 2000; Silk, 1997].

A unique role of the universe is in creating the environment in which galaxies, stars, and planets develop, thus providing a setting in which local physics and chemistry can function in a way that enables the evolution of life on planets such as the Earth. If the cosmological environment were substantially different, local conditions would be different and in most cases we would not be here [Carr and Rees, 1979; Davies, 1982; Barrow and Tipler, 1984; Tegmark, 1998; Rees, 1999] — indeed no biological evolution at all would have taken place. Thus cosmology is of substantial interest to the whole of the scientific endeavor, for it sets the

framework for the rest of science, and indeed for the very existence of observers and scientists. It is unique as the ultimate historical/geographical science.

Cosmology as a serious scientific study began with the discovery of Einstein's static universe in 1917, followed by the key observational discovery of the linear redshift-distance relation by Hubble in 1929, indicating the expansion of the universe, and the development of theoretical understanding of the geometry and dynamics of the non-static Friedmann-Lemaître models with their Robertson-Walker geometry [North, 1965; Berendzen *et al.*, 1976; Smith, 1982; Ellis, 1989; Kragh, 1996]. It has been transformed in the past decades into a mainstream branch of physics [Barnett *et al.*, 1996; Nilsson *et al.*, 1991] by the linking of nuclear and particle physics theories to observable features of the cosmos [Weinberg, 1972; Kolb and Turner, 1990; Peacock, 1999; Allday, 2002; Dodelson, 2003], and into an important part of astronomy because of the massive flow of new astronomical data becoming available [Gunn *et al.*, 1978; Harwit, 1984; Bothun, 1998], particularly through new ground-based telescopes such as Keck and through balloon and satellite observatories such as the Hubble Space telescope (optical and ultraviolet), IRAS (infra-red), ROSAT (x-ray), and COBE and WMAP (microwave). Thus the subject has progressed from a mainly mathematical and even philosophical exercise to an important part of mainstream science, with a well-established standard model confirmed by various strands of evidence [Weinberg, 1972; Peebles *et al.*, 1991; Silk, 1997; Peacock, 1999; Dodelson, 2003]. Nevertheless because of its nature, it is different from any other branch of the natural sciences, its unique features playing themselves out in the ongoing interaction between speculation, theory, and observation.

Cosmology's major difference from the other sciences is the uniqueness of its object of study — the Universe as a whole [McCrea, 1953; McCrea, 1960; Munitz, 1962] — together with its role as the background for all the rest of physics and science, the resulting problems being accentuated by the vast scale of the universe and by the extreme energies occurring in the very early universe. We are unable to manipulate in any way its originating conditions, and there are limitations on our ability to observe both to very distant regions and to very early times. Additionally, there are limits to our ability to test the physics relevant at the earliest epochs. Consequently it is inevitable that (as is also the case for the other historical sciences) philosophical choices will to some degree shape the nature of cosmological theory, particularly when it moves beyond the purely descriptive to an explanatory role [Matravers *et al.*, 1995] — that move being central to its impressive progress in recent decades. These philosophical choices will strongly influence the resulting understanding, and even more so if we pursue a theory with more ambitious explanatory aims.

After a substantial outline of present day cosmology in Section 2, these issues will be explored in the subsequent sections, based on a series of thirty-four *Theses* clustered around nine key aspects of the nature of cosmology, broadly speaking relating to geometry, physics, and philosophy, that frame the context of the philosophical issues facing cosmology and its relation to local physics. I believe this

formulation helps focus on specific issues of importance in this relation. To those who believe cosmology is simply about determining a number of physical parameters, this will seem a vastly over-complicated approach; but then a major purpose of this paper is precisely to counter such simplistic visions of the nature of cosmology. For other reports on the philosophy of cosmology, see [McCrea, 1970; Munitz, 1962; Ellis, 1991; Leslie, 1994; Leslie, 1998].

2 OUTLINE OF COSMOLOGY

A series of basic features of present day cosmology are now well established. Decades of painstaking work has established the distances of galaxies and hence the huge scale of the universe, as well as the basic feature that the universe is expanding and evolving; the old dream of a static universe is unviable [Ellis, 1990]. Cosmology proceeds by assuming *the laws of physics are the same everywhere, and underlie the evolution of the universe*. The dominant role of gravity, despite its weakness, then arises from the fact that it is the only known force acting effectively on astronomical scales (the other known long-range force is electromagnetism, but in this case negative charges balance out positive charges, leaving no resultant large-scale effect). Consequently, cosmological theory describing all but the very earliest times is based on the classical relativistic theory of gravitation, namely Einstein's General Theory of Relativity [Malament, 2006], with the matter present determining space-time curvature and hence the evolution of the universe. The way this works out in any particular situation depends on the nature of the matter/fields present, described by their effective equations of state and interaction potentials.

The survey of cosmology in this section looks successively at the basic models of cosmology; the hot big bang; cosmological observations, including the Cosmic Background Radiation anisotropy spectrum; causal and visual horizons, and their implications; recent theoretical developments (including inflation); the very early universe; and the present concordance model, which includes both dark matter and dark energy.

2.1 Basic Theory

Cosmology starts by assuming that *the large-scale evolution of spacetime can be determined by applying Einstein's field equations of Gravitation ('EFE') everywhere*: global evolution will follow from local physics. The standard models of cosmology [Robertson, 1933; Ehlers, 1993; Weinberg, 1972; Hawking and Ellis, 1973] are based on the assumption that once one has averaged over a large enough physical scale, *isotropy is observed by all fundamental observers* (the preferred family of observers associated with the average motion of matter in the universe). When this isotropy is exact, *the universe is spatially homogeneous as well as isotropic* [Walker, 1944; Ehlers, 1993; Ellis, 1971a]. The matter motion is then along irrotational and shear-free geodesic curves with tangent vector u^a , implying the existence of a canoni-

cal time-variable t obeying $u_a = -t_{,a}$. The Robertson-Walker ('RW') geometries used to describe the large-scale structure of the universe [Robertson, 1935; Walker, 1936] embody these symmetries exactly. Consequently they are conformally flat, that is, the Weyl tensor is zero:

$$(1) \quad C_{ijkl} := R_{ijkl} + \frac{1}{2}(R_{ik}g_{jl} + R_{jl}g_{ik} - R_{il}g_{jk} - R_{jk}g_{il}) - \frac{1}{6}R(g_{ik}g_{jl} - g_{il}g_{jk}) = 0;$$

this tensor represents the free gravitational field, enabling non-local effects such as tidal forces and gravitational waves which do not occur in the exact RW geometries.

Comoving coordinates can be chosen so that the metric takes the form:

$$(2) \quad ds^2 = -dt^2 + S^2(t) d\sigma^2, \quad u^a = \delta^a_0 \quad (a = 0, 1, 2, 3)$$

where $S(t)$ is the time-dependent scale factor, and the worldlines with tangent vector $u^a = dx^a/dt$ represent the histories of fundamental observers. The space sections $\{t = const\}$ are surfaces of homogeneity and have maximal symmetry: they are 3-spaces of constant curvature $K = k/S^2(t)$ where k is the sign of K . The normalized metric $d\sigma^2$ characterizes a 3-space of normalized constant curvature k ; coordinates (r, θ, ϕ) can be chosen such that

$$(3) \quad d\sigma^2 = dr^2 + f^2(r) (d\theta^2 + \sin^2 \theta d\phi^2)$$

where $f(r) = \{\sin r, r, \sinh r\}$ if $k = \{+1, 0, -1\}$ respectively. The rate of expansion at any time t is characterised by the *Hubble parameter* $H(t) = \dot{S}/S$.

To determine the metric's evolution in time, one applies the Einstein Field Equations ('EFE'), showing the effect of matter on space-time curvature, to the metric (2,3). Because of local isotropy, the matter tensor T_{ab} necessarily takes a perfect fluid form relative to the preferred worldlines with tangent vector u^a :

$$(4) \quad T_{ab} = (\mu + p/c^2)u_a u_b + (p/c^2)g_{ab}$$

(c is the speed of light). The energy density $\mu(t)$ and pressure term $p(t)/c^2$ are the timelike and spacelike eigenvalues of T_{ab} . The integrability conditions for the EFE are the *energy-density conservation equation*

$$(5) \quad T^{ab}_{;b} = 0 \Leftrightarrow \dot{\mu} + (\mu + p/c^2)3\dot{S}/S = 0.$$

This becomes determinate when a suitable equation of state function $w := pc^2/\mu$ relates the pressure p to the energy density μ and temperature T : $p = w(\mu, T)\mu/c^2$ (w may or may not be constant). Baryons have $\{p_{bar} = 0 \Leftrightarrow w = 0\}$ and radiation has $\{p_{rad}c^2 = \mu_{rad}/3 \Leftrightarrow w = 1/3, \mu_{rad} = aT_{rad}^4\}$, which by (5) imply

$$(6) \quad \mu_{bar} \propto S^{-3}, \quad \mu_{rad} \propto S^{-4}, \quad T_{rad} \propto S^{-1}.$$

The scale factor $S(t)$ obeys the *Raychaudhuri equation*

$$(7) \quad 3\ddot{S}/S = -\frac{1}{2}\kappa(\mu + 3p/c^2) + \Lambda,$$

where κ is the gravitational constant and Λ the cosmological constant.¹ This shows that the active gravitational mass density of the matter and fields present is $\mu_{grav} := \mu + 3p/c^2$. For ordinary matter this will be positive:

$$(8) \quad \mu + 3p/c^2 > 0 \Leftrightarrow w > -1/3$$

(the ‘Strong Energy Condition’), so ordinary matter will tend to cause the universe to decelerate ($\dot{S} < 0$). It is also apparent that a positive cosmological constant on its own will cause an accelerating expansion ($\dot{S} > 0$). When matter and a cosmological constant are both present, either result may occur depending on which effect is dominant. The first integral of equations (5, 7) when $\dot{S} \neq 0$ is the *Friedmann equation*

$$(9) \quad \frac{\dot{S}^2}{S^2} = \frac{\kappa\mu}{3} + \frac{\Lambda}{3} - \frac{k}{S^2}.$$

This is just the Gauss equation relating the 3-space curvature to the 4-space curvature, showing how matter directly causes a curvature of 3-spaces [Ehlers, 1993; Ellis, 1971a]. Because of the spacetime symmetries, the ten EFE are equivalent to the two equations (7, 9). Models of this kind, that is with a Robertson-Walker (‘RW’) geometry with metric (2, 3) and dynamics governed by equations (5, 7, 9), are called *Friedmann-Lemaître universes* (‘FL’ for short). The Friedmann equation (9) controls the expansion of the universe, and the conservation equation (5) controls the density of matter as the universe expands; when $\dot{S} \neq 0$, equation (7) will necessarily hold if (5, 9) are both satisfied.

Given a determinate matter description (specifying the equation of state $w = w(\mu, T)$ explicitly or implicitly) for each matter component, the existence and uniqueness of solutions follows both for a single matter component and for a combination of different kinds of matter, for example $\mu = \mu_{bar} + \mu_{rad} + \mu_{cdm} + \mu_{\nu}$ where we include cold dark matter (cdm) and neutrinos (ν). Initial data for such solutions at an arbitrary time t_0 (eg. today) consists of,

- The *Hubble constant* $H_0 := (\dot{S}/S)_0 = 100h$ km/sec/Mpc;
- A dimensionless *density parameter* $\Omega_{i0} := \kappa\mu_{i0}/3H_0^2$ for each type of matter present (labelled by i);
- If $\Lambda \neq 0$, either $\Omega_{\Lambda 0} := \Lambda/3H_0^2$, or the dimensionless *deceleration parameter* $q_0 := -(\ddot{S}/S)_0 H_0^{-2}$.

Given the equations of state for the matter, this data then determines a unique solution $\{S(t), \mu(t)\}$, i.e. a unique corresponding universe history. The total matter density is the sum of the terms Ω_{i0} for each type of matter present, for example

$$(10) \quad \Omega_{m0} = \Omega_{bar0} + \Omega_{rad0} + \Omega_{cdm0} + \Omega_{\nu 0},$$

¹A cosmological constant can also be regarded as a fluid with pressure p related to the energy density μ by $\{p = -\mu c^2 \Leftrightarrow w = -1\}$. For the history of the cosmological constant, see [Earman, 2001; Earman, 2003].

and the total density parameter Ω_0 is the sum of that for matter and for the cosmological constant:

$$(11) \quad \Omega_0 = \Omega_{m0} + \Omega_{\Lambda 0}.$$

Evaluating the Raychaudhuri equation (7) at the present time gives an important relation between these parameters: when the pressure term p/c^2 can be ignored relative to the matter term μ (as is plausible at the present time),²

$$(12) \quad q_0 = \frac{1}{2} \Omega_{m0} - \Omega_{\Lambda 0}.$$

This shows that a cosmological constant Λ can cause an acceleration (negative q_0); if it vanishes, the expression simplifies: $\Lambda = 0 \Rightarrow q = \frac{1}{2} \Omega_{m0}$, showing how matter causes a deceleration of the universe. Evaluating the Friedmann equation (9) at the time t_0 , the spatial curvature is

$$(13) \quad K_0 := k/S_0^2 = H_0^2 (\Omega_0 - 1).$$

The value $\Omega_0 = 1$ corresponds to spatially flat universes ($K_0 = 0$), separating models with positive spatial curvature ($\Omega_0 > 1 \Leftrightarrow K_0 > 0$) from those with negative spatial curvature ($\Omega_0 < 1 \Leftrightarrow K_0 < 0$).

The FL models are the standard models of modern cosmology, surprisingly effective in view of their extreme geometrical simplicity. One of their great strengths is their explanatory role in terms of making explicit the way the local gravitational effect of matter and radiation determines the evolution of the universe as a whole, this in turn forming the dynamic background for local physics (including the evolution of the matter and radiation).

2.1.1 The basic solutions

For baryons (pressure-free matter) and non-interacting radiation, the Friedmann equation (9) takes the form

$$(14) \quad \frac{3\dot{S}^2}{S^2} = \frac{A}{S^3} + \frac{B}{S^4} + \frac{\Lambda}{3} - \frac{3k}{S^2}$$

where $A := \kappa\mu_{bar0}S_0^3$ and $B := \kappa\mu_{rad0}S_0^4$. The behaviour depends on the cosmological constant Λ [Robertson, 1933; Rindler, 2001].

When $\Lambda = 0$, the universe starts off at a very dense initial state — according to the classical theory, an initial singularity where the density and curvature go infinite (see Sec. 2.1.2). Its future fate depends on the value of the spatial curvature, or equivalently the density parameter Ω_0 . The universe expands forever if $\{k = 0 \Leftrightarrow \Omega_0 = 1\}$ or $\{k < 0 \Leftrightarrow \Omega_0 < 1\}$, but collapses to a future singularity if $\{k > 0 \Leftrightarrow \Omega_0 > 1\}$. Thus $\Omega_0 = 1$ corresponds to the critical density μ_{crit} separating $\Lambda = 0$ FL models that recollapse in the future from those that expand forever, and Ω_0 is just the ratio of the matter density to this critical density:

²Assuming we represent ‘dark energy’ (Sec. 2.3.6) as a cosmological constant.

$$(15) \{ \Omega_{crit} = 1 \Leftrightarrow \kappa\mu_{crit} = 3H_0^2 \} \Rightarrow \Omega_0 := \kappa\mu_0/3H_0^2 = \mu_0/\mu_{crit} .$$

When $\Lambda < 0$, all solutions start at a singularity and recollapse.

When $\Lambda > 0$, if $k = 0$ or $k = -1$ all solutions start at a singularity and expand forever. If $k = +1$ there can again be models with a singular start, either expanding forever or collapsing to a future singularity. However in this case a static solution (the Einstein static universe) is also possible, as well as models asymptotic to this static state in either the future or the past. Furthermore models with $k = +1$ can bounce (collapsing from infinity to a minimum radius and re-expanding).

The dynamical behaviour of these models has been investigated in depth: first for dust plus a cosmological constant [Robertson, 1933; Rindler, 2001], followed by perfect fluids, fluids with bulk viscosity, kinetic theory solutions, and scalar field solutions. Current models employ a realistic mixture of matter components (baryons, radiation, neutrinos, cold dark matter, a scalar field, and perhaps a cosmological constant). Informative phase planes show clearly the way higher symmetry (self-similar) models act as attractors and saddle points for the other models [Madsen and Ellis, 1988; Ehlers and Rindler, 1989].

The simplest expanding solutions are the following:

1. The *Einstein-de Sitter* model, for which $\{p = 0, \Lambda = 0, k = 0\} \Rightarrow \Omega_0 = 1$. This is the simplest expanding non-empty solution:

$$(16) S(t) = C t^{2/3}$$

starting from a singular state at time $t = 0$ (C is an arbitrary constant). Its age (the proper time since the start of the universe) when the Hubble constant takes the value H_0 is $\tau_0 = \frac{2}{3H_0}$. This is a good model of the expansion of the universe since radiation domination ended until the recent times when a cosmological constant started to dominate the expansion. It is also a good model of the far future universe if $k = 0$ and $\Lambda = 0$.

2. The *Milne* model, for which $\{\mu = p = 0, \Lambda = 0, k = -1\} \Rightarrow \Omega_0 = 0$, giving a linearly expanding empty solution:

$$(17) S(t) = C t.$$

This is just flat space-time as seen by a uniformly expanding set of observers [Rindler, 2001, pp. 360-363], singular at $t = 0$. Its age is $\tau_0 = \frac{1}{H_0}$. It is a good model of the far future universe if $k < 0$ and $\Lambda = 0$.

3. The *de Sitter* universe, for which $\{\mu = p = 0, \Lambda \neq 0, k = 0\} \Rightarrow \Omega_0 = 0$, giving the steady state expanding empty solution:³

$$(18) S(t) = C \exp(Ht),$$

³The Steady State universe of Bondi, Hold and Hoyle [Bondi, 1960] utilised this metric, but was non-empty as they abandoned the EFE.

where C and H are constants. As the expansion rate is constant forever, there is no start and its age is infinite.⁴ It is a good model of the far future universe for those cases which expand forever with $\Lambda > 0$. It can alternatively be understood as a solution with $\Lambda = 0$ and containing matter with the exceptional equation of state $\mu + p/c^2 = 0$. There are other RW forms of the de Sitter Universe: a geodesically complete form with $k = +1$, $S(t) = S_0 \cosh Ht$ (a regular bounce), and another geodesically incomplete form with $k = -1$, $S(t) = S_0 \sinh Ht$ (a singular start). This lack of uniqueness is possible because *this is a spacetime of constant curvature*, with no preferred timelike directions or space sections [Schrödinger, 1956; Hawking and Ellis, 1973; Rindler, 2001].⁵

2.1.2 An initial singularity?

The above are specific models: what can one say generically? When the inequality (8) is satisfied, one obtains directly from the Raychaudhuri equation (7) the

Friedmann-Lemaître Universe Singularity Theorem [Ehlers, 1993; Ellis, 1971a]: In a FL universe with $\Lambda \leq 0$ and $\mu + 3p/c^2 > 0$ at all times, at any instant t_0 when $H_0 \equiv (\dot{S}/S)_0 > 0$ there is a finite time t_* : $t_0 - (1/H_0) < t_* < t_0$, such that $S(t) \rightarrow 0$ as $t \rightarrow t_*$; the universe starts at a space-time singularity there, with $\mu \rightarrow \infty$ and $T \rightarrow \infty$ if $\mu + p/c^2 > 0$.

This is not merely a start to matter — it is a start to space, to time, to physics itself. It is the most dramatic event in the history of the universe: it is the start of existence of everything. The underlying physical feature is the non-linear nature of the EFE: going back into the past, the more the universe contracts, the higher the active gravitational density, causing it to contract even more. The pressure p that one might have hoped would help stave off the collapse makes it even worse because (consequent on the form of the EFE) p enters algebraically into the Raychaudhuri equation (7) with the same sign as the energy density μ . Note that the Hubble constant gives an estimate of the age of the universe: the time $\tau_0 = t_0 - t_*$ since the start of the universe is less than $1/H_0$.

This conclusion can in principle be avoided by a cosmological constant, but in practice this cannot work because we know the universe has expanded by at least a ratio of 6, as we have seen objects at a redshift⁶ of 5; from (14), the cosmological constant would have to have an effective magnitude at least $6^3 = 216$ times the present matter density to dominate and cause a turn-around then or at any earlier time, and so would be much bigger than its observed present upper limit (of the

⁴It is however singular in that it is geodesically incomplete; this metric covers only half the de Sitter hyperboloid [Schrödinger, 1956; Hawking and Ellis, 1973].

⁵There is also a static (non-RW) form of the metric — the first form of the metric discovered.

⁶The redshift z for light emitted at t_e and observed at t_0 is related to the expansion by $1 + z = S(t_0)/S(t_e)$, see Sec. 2.3.3.

same order as the present matter density). Accordingly, no turn around is possible while classical physics holds [Ehlers and Rindler, 1989]. However energy-violating matter components such as a scalar field (Sec. 2.6) can avoid this conclusion, if they dominate at early enough times; but this can only be when quantum fields are significant, when the universe was at least 10^{12} smaller than at present.

Because $T_{rad} \propto S^{-1}$ (eqn.(6)), a major conclusion is that *a Hot Big Bang must have occurred; densities and temperatures must have risen at least to high enough energies that quantum fields were significant*, at something like the GUT energy. The universe must have reached those extreme temperatures and energies at which classical theory breaks down.

2.2 The hot big bang

The matter and radiation in the universe gets hotter and hotter as we go back in time towards the initial quantum state, because it was compressed into a smaller volume. In this *Hot Big Bang* epoch in the early universe, we can use standard physical laws to examine the processes going on in the expanding mixture of matter and radiation [Weinberg, 1972; Perkins, 2005]. A key feature is that about 300,000 years after the start of the Hot Big Bang epoch, nuclei and electrons combined to form atoms. At earlier times when the temperature was higher, atoms could not exist, as the radiation then had so much energy it disrupted any atoms that tried to form into their constituent parts (nuclei and electrons). Thus at earlier times matter was ionized, consisting of negatively charged electrons moving independently of positively charged atomic nuclei. Under these conditions, the free electrons interact strongly with radiation by Thomson scattering. Consequently matter and radiation were tightly coupled in equilibrium at those times, and the Universe was opaque to radiation. When the temperature dropped through the ionization temperature of about 4000K, atoms formed from the nuclei and electrons, and this scattering ceased: the Universe became very transparent (today we are able to see galaxies at enormous distances from us). The time when this transition took place is known as the *time of decoupling* — it was the time when matter and radiation ceased to be tightly coupled to each other, at a redshift $z_{dec} \simeq 1100$ [Dodelson, 2003]. By (6), the universe was radiation dominated ($\mu_{rad} \gg \mu_{mat}$) at early times and matter dominated ($\mu_{rad} \ll \mu_{mat}$) at late times;⁷ matter-radiation density equality occurred significantly before decoupling (the temperature T_{eq} when this equality occurred was $T_{eq} \simeq 10^4\text{K}$; at that time the scale factor was $S_{eq} \simeq 10^4 S_0$, where S_0 is the present-day value). The dynamics of both the background model and of perturbations about that model differ significantly before and after S_{eq} [Dodelson, 2003].

⁷The dynamically dominant Cold Dark Matter (Sec. 2.3.6) obeys the same density law (6) as baryons.

2.2.1 Cosmic Blackbody Radiation

Radiation was emitted by matter at the time of decoupling, thereafter travelling freely to us through the intervening space. When it was emitted, it had the form of blackbody radiation, because this is a consequence of matter and radiation being in thermodynamic equilibrium at earlier times. Thus *the matter at $z = z_{dec}$ forms the Last Scattering Surface (LSS) in the early universe, emitting Cosmic Blackbody Background Radiation*⁸ ('CBBR') at 4000K, that since then has travelled freely with its temperature T scaling inversely with the scale function of the universe.⁹ As the radiation travelled towards us, the universe expanded by a factor of about 1100; consequently by the time it reaches us, it has cooled to 2.75 K (that is, about 3 degrees above absolute zero, with a spectrum peaking in the microwave region), and so is extremely hard to observe. It was however detected in 1965, and its spectrum has since been intensively investigated, its blackbody nature being confirmed to high accuracy [Partridge, 1995]. Its existence is now taken as solid proof both that the Universe has indeed expanded from a hot early phase, and that standard physics applied unchanged at that era in the early universe.

2.2.2 Particle interactions and element formation

The thermal capacity of the radiation is hugely greater than that of the matter. At very early times before decoupling, the temperatures of the matter and radiation were the same (because they were in equilibrium with each other), scaling as $1/S(t)$ (eqn.(6)). The early universe exceeded any temperature that can ever be attained on Earth or even in the centre of the Sun; as it dropped towards its present value of 3 K, successive physical reactions took place that determined the nature of the matter we see around us today. At very early times and high temperatures, only elementary particles can survive and even neutrinos had a very small mean free path; as the universe cooled down, neutrinos decoupled from the matter and streamed freely through space. At these times the expansion of the universe was radiation dominated, and we can approximate the universe then by models with $\{k = 0, w = 1/3, \Lambda = 0\}$, the resulting simple solution of (14) uniquely relating time to temperature:

$$(19) \quad S(t) = S_0 t^{1/2}, \quad t = 1.92 \text{ sec} \left[\frac{T}{10^{10} \text{ K}} \right]^{-2}.$$

(There are no free constants in the latter equation).

At very early times, even neutrinos were tightly coupled and in equilibrium with the radiation; they decoupled at about 10^{10} K [Dodelson, 2003, pp. 44-46], resulting in a relic neutrino background density in the universe today of about $\Omega_{\nu 0} \simeq 10^{-5}$ if they are massless (but it could be higher depending on their masses). Key events in the early universe are associated with out of equilibrium phenomena

⁸This is often called "Cosmic Microwave Background", or CMB for short. However it is only microwave at the present epoch.

⁹This scaling for freely propagating radiation follows from the discussion in Sec. 2.3.3.

[Dodelson, 2003, p. 58]. An important event was the era of *nucleosynthesis*, the time when the light elements were formed. Above about 10^9K , nuclei could not exist because the radiation was so energetic that as fast as they formed, they were disrupted into their constituent parts (protons and neutrons). However below this temperature, if particles collided with each other with sufficient energy for nuclear reactions to take place, the resultant nuclei remained intact (the radiation being less energetic than their binding energy and hence unable to disrupt them). Thus the nuclei of the light elements — deuterium, tritium, helium, and lithium — were created by neutron capture. This process ceased when the temperature dropped below about 10^8 K (the nuclear reaction threshold). In this way, the proportions of these light elements at the end of nucleosynthesis were determined; they have remained virtually unchanged since. The rate of reaction was extremely high; all this took place within the first three minutes of the expansion of the Universe. One of the major triumphs of Big Bang theory is that *theory and observation are in excellent agreement provided the density of baryons is low: $\Omega_{\text{bar}0} \simeq 0.044$. Then the predicted abundances of these elements (25% Helium by weight, 75% Hydrogen, the others being less than 1%) agrees very closely with the observed abundances.* Thus the standard model explains the origin of the light elements in terms of known nuclear reactions taking place in the early Universe [Schramm and Turner, 1998]. However heavier elements cannot form in the time available (about 3 minutes).

In a similar way, physical processes in the very early Universe (before nucleosynthesis) can be invoked to explain the ratio of matter to anti-matter in the present-day Universe: a small excess of matter over anti-matter must be created then in the process of *baryosynthesis*, without which we could not exist today (if there were no such excess, matter and antimatter would have all annihilated to give just radiation [Silk, 2005]). However other quantities (such as electric charge) are believed to have been conserved even in the extreme conditions of the early Universe, so their present values result from given initial conditions at the origin of the Universe, rather than from physical processes taking place as it evolved. In the case of electric charge, the total conserved quantity appears to be zero: after quarks form protons and neutrons at the time of baryosynthesis, there are equal numbers of positively charged protons and negatively charged electrons, so that at the time of decoupling there were just enough electrons to combine with the nuclei and form uncharged atoms (it seems there is no net electrical charge on astronomical bodies such as our galaxy; were this not true, electromagnetic forces would dominate cosmology, rather than gravity).

After decoupling, matter formed large scale structures through gravitational instability [Bothun, 1998, pp. 183-222] which eventually led to the formation of the first generation of stars [Silk, 2005] and is probably associated with the re-ionization of matter [Dodelson, 2003, p. 73]. However at that time planets could not form for a very important reason: there were no heavy elements present in the Universe. The first stars aggregated matter together by gravitational attraction, the matter heating up as it became more and more concentrated, until its temperature exceeded the thermonuclear ignition point and nuclear reactions started

burning hydrogen to form helium. Eventually more complex nuclear reactions started in concentric spheres around the centre, leading to a build-up of heavy elements (carbon, nitrogen, oxygen for example), up to iron. These elements can form in stars because there is a long time available (millions of years) for the reactions to take place. Massive stars burn relatively rapidly, and eventually run out of nuclear fuel. The star becomes unstable, and its core rapidly collapses because of gravitational attraction. The consequent rise in temperature blows it apart in a giant explosion, during which time new reactions take place that generate elements heavier than iron; this explosion is seen by us as a Supernova (“New Star”) suddenly blazing in the sky, where previously there was just an ordinary star. Such explosions blow into space the heavy elements that had been accumulating in the star’s interior, forming vast filaments of dust around the remnant of the star. It is this material that can later be accumulated, during the formation of second generation stars, to form planetary systems around those stars. Thus *the elements of which we are made (the carbon, nitrogen, oxygen and iron nuclei for example) were created in the extreme heat of stellar interiors, and made available for our use by supernova explosions.* Without these explosions, we could not exist.

2.3 Cosmological Observations

Cosmological models only become meaningful when related to astronomical observations [Hoyle, 1960; Sandage, 1961; Ellis, 1971a; Weinberg, 1972]. These are of two kinds: astronomical observations of distant matter tells us what was happening far away in the universe and (because of the finite speed of light) a long time ago. On the other hand observations of nearby objects (matter on Earth, the solar system, nearby stars for example) when related to theories of origins tell us what was happening very near our past world line a very long time ago. The first set of observations may be characterized as “null cone” observations, the second as “geological” observations, one of the most important being the determination of local element abundances, which are then related to nucleosynthesis calculations (Sec. 2.2.2).

Observations are totally dependent on telescope and detector technology [Harwit, 1984; Bothun, 1998]. After the initial establishment of distance scales and the basic evidence of cosmic homogeneity and expansion in the 1920s and 1930s, progress was slow until the 1960s when observations were extended from the optical to the entire electromagnetic spectrum. In recent decades cosmology has changed from a data-poor to a data-rich subject. Massive new data sets are now available because of the extraordinary improvement of telescope, detector, and computer technology in recent decades, particularly the advent of new detectors such as Charge Coupled Devices (CCD’s) and fibre optics (enabling simultaneous measurement of hundreds of redshifts). We now have not only optical, ultraviolet, and infrared observations of galaxies, determining luminosities and spectra with unprecedented sensitivity, but also radio, X-ray, and gamma-ray sky surveys. Galaxies have been detected up to a redshift of 6 and we have identified

many quasi-stellar objects and gamma-ray bursters as well as multiple images of very distant gravitationally-lensed galaxies [Harwit, 1984]. Large-scale structures (clusters of galaxies, superclusters, walls, and voids) have been identified, with associated large-scale velocity flows [Bothun, 1998, pp. 85-137].

In addition to large-scale number-count and redshift surveys, we have measured the background radiation spectrum and anisotropies at all wavelengths. We identify the radiation as ‘background’ precisely when it is constant on very small angular scales (as opposed to discrete sources, which appear as isolated objects). There is a complex relation of this radiation to the intergalactic matter density and thermal history. The most important component of the background radiation is the Cosmic Blackbody Radiation (‘CBR’) mentioned above (Sec. 2.2); detailed observations have mapped its temperature over the whole sky at a sensitivity of better than one part in 10^5 . However other components of the background radiation (X-ray and radio in particular) convey important information on the temperature and density of intergalactic matter, and hence strongly restrict its possible thermal history. For example hot matter emits X-rays, so the X-ray background measurement restricts the amount of hot intergalactic matter allowed; while neutral hydrogen strongly absorbs Ultra-Violet radiation to give the Lyman alpha spectral absorption line, so absence of such absorption gives strong limits on the amount of neutral hydrogen and hence on the temperature of intergalactic matter.

2.3.1 Isotropy

The first important point about cosmological observations is that *when averaged on a large enough physical scale (clusters of galaxies and above) they are statistically isotropic about us*; there is no direction apparently pointing to the centre of the universe. The *high degree of isotropy of the CBR* strongly supports this conclusion: its temperature is the same in all directions about us to better than one part in 10,000 after we have allowed for the motion of the Earth relative to the cosmos (about 250 km/sec), which creates a temperature dipole at one part in a thousand.¹⁰ Any inhomogeneities or anisotropies in the matter distribution lead to anisotropies in this radiation, as recently measured at only one part in 10^5 by the extremely sensitive detectors of the COBE and WMAP satellites. This high degree of isotropy is the major reason we believe the Universe itself is spatially homogeneous and isotropic to a good approximation (see Sec. 4.2.2), providing good observational support for use of the FL universe models as accurate models of the observed region of the universe.

2.3.2 Distance scale and ages

The underlying problem in all astronomy is determining the distances of observed objects. This is done by a ‘cosmic distance ladder’ [Bothun, 1998, pp. 25-83]

¹⁰The CBR dipole that could be interpreted as due to a major cosmological inhomogeneity is rather interpreted as being due to our motion (‘peculiar velocity’) relative to a spatially homogeneous universe.

whereby nearest objects have their distance determined by parallax (i.e. essentially by local trigonometry); and more distant ones by a series of consecutive distance indicators (Cepheid variables, RR Lyrae variables, brightest red supergiants) until at a cosmological distance, redshift z is a primary distance indicator, but is contaminated by local velocities of matter relative to the rest-frame of the universe. Other distance indicators (for example the Tully-Fisher method, the luminosity function of planetary nebulae, the globular cluster luminosity function, surface brightness fluctuations) serve to refine the estimates [Bothun, 1998].

Closely associated with the distance scale is determination of the Hubble constant H_0 (the present rate of expansion of the universe), because estimates of the size of the observable region of the universe scale with the Hubble constant. But the Hubble constant also determines the age of the universe, so its determination underlies a crucial consistency condition for cosmology: *the age of objects in the universe (rocks, planets, stars, star clusters, galaxies) must be less than the age of the universe*. This condition has been a cause of concern ever since we have had good estimates of ages and of the Hubble constant.¹¹ It seems not to be violated by current observations of low redshift objects given the current estimates of $H_0 \simeq 70$ km/sec/Mpc, giving an age of the universe of about 15 billion years whereas the oldest star clusters seem to be about 14 billion years old. However it is very tight, perhaps even problematic, for very distant (and so much younger) objects [Jain and Dev, 2005].

2.3.3 Observational relations

Light travels on null geodesics $x^a(\lambda)$ in spacetime (the tangent vector $k^a := dx^a/d\lambda$ is such that $k^a{}_{;b}k^b = 0$, $k^ak_a = 0$). In a RW geometry, it suffices to consider only radial null geodesics (by the symmetries of the model, these are equivalent to generic geodesics). Then from (2) we find that for light emitted at time t_e and received at time t_0 , the comoving radial distance $u(t_0, t_1) := r_0 - r_1$ between comoving emitters and receivers is given by

$$(20) \quad \{ds^2 = 0, d\theta = 0 = d\phi\} \Rightarrow u(t_0, t_1) = \int_{t_1}^{t_0} \frac{dt}{S(t)} = \int_{S_1}^{S_0} \frac{dS}{S\dot{S}}$$

with \dot{S} given by the Friedmann equation (9). The key quantities related to cosmological observations are redshift, area distance (or apparent size), and local volume corresponding to some increment in distance (determining number counts) [Sandage, 1961; Ellis, 1971a; Weinberg, 1972]. The redshift z measured for an object emitting light of wavelength λ_e that is observed with wavelength λ_0 is given by

¹¹Indeed Hubble himself never fully accepted the expanding universe theory because of age difficulties, preferring to refer to a redshift-distance relation rather than a velocity distance relation [Hubble, 1936]. However the problem has been eased by a series of revisions of the value of the Hubble constant since then, due to a better understanding of the primary distance indicators.

$$(21) \quad 1 + z := \frac{\lambda_0}{\lambda_e} = (1 + z_c)(1 + z_v),$$

where z_v is the redshift caused by the local peculiar motion of the object observed ($z_v = 0$ for comoving objects), and z_c is the cosmological redshift given by

$$(22) \quad 1 + z_c = \frac{S(t_0)}{S(t_e)}.$$

From eqn.(21), the same ratio of observed to emitted light holds for all wavelengths: a key identifying property of redshift. The problem in using redshifts of objects as a distance indicator is to separate out the cosmological from the Doppler components, which lead to redshift-space distortions [Dodelson, 2003, pp. 275-282]; this can reasonably be done for a cluster of galaxies by appropriate averaging over cluster members ($\langle z_v \rangle = 0$ for a comoving cluster). The area distance r_0 of an object at redshift z_c and of linear size l which subtends angular size α is given by¹²

$$(23) \quad r_0(z_c) := \frac{l}{\alpha} = f(u)S(t_e).$$

Thus measures of apparent sizes will determine the area distance if the source physical size is known. The flux of radiation F measured from a point source of luminosity L emitting radiation isotropically is given by the fraction of radiant energy emitted by the source in a unit of time that is received by the telescope:

$$(24) \quad F = \frac{L}{4\pi} \frac{1}{f^2(u)S^2(t_0)(1+z)^2};$$

(the two redshift factors account firstly for time dilation observed between observer and source, and secondly for loss of energy due to redshifting of photons). The source's *apparent magnitude* m is defined from the flux: $m = -2.5\log_{10}F + const.$ On using (22, 23), equation (24) becomes

$$(25) \quad F = \frac{L}{4\pi} \frac{1}{r_0^2(1+z)^4}.$$

showing that measures of magnitudes will determine the area distance if the source's intrinsic luminosity is known. On using (23) it follows from (25) that *the point-wise surface brightness of extended objects (the flux received per unit solid angle) depends only on redshift* [Kristian and Sachs, 1966; Ellis, 1971a] — a key feature in determining detection probabilities and in gravitational lensing observations. It further follows from this result that *a blackbody spectrum emitted at temperature T_e when observed with a redshift z remains a blackbody spectrum but with observed temperature $T_0 = T_e/(1+z)$* — a crucial feature in analyzing the CBR observations.

¹²This depends only on z_c because apparent shapes and sizes are independent of the motion of the source.

Using the Friedmann equation and the relevant equation of state for matter, the area distance can be determined as a function of redshift z_c in terms of the Hubble constant H_0 , deceleration parameter q_0 , and cosmological constant Λ . In the case of pressure-free matter with vanishing cosmological constant, one obtains from (20), (9), (22), and (23)¹³ the Mattig relation [Sandage, 1961]

$$(26) \quad r_0(z_c) = \frac{1}{H_0 q_0^2} \frac{(q_0 - 1)(1 + 2q_0 z_c)^{1/2} + (q_0(z_c - 1) + 1)}{(1 + z_c)^2}.$$

Consequently measures of either apparent size of sources of known physical size, or of radiant flux from sources of known intrinsic luminosity, will determine the deceleration parameter q_0 . Generalizations of this relation hold if a cosmological constant or radiation is present. An interesting aspect is that there is a minimum apparent size for objects of fixed physical size at some redshift $z_c = z_*$ depending on the density parameter and the cosmological constant. The past light cone of the observer attains a maximum area at z_* ; the entire universe acts as a gravitational lens for further off objects, magnifying their apparent size so that very distant objects can appear to have the same angular size as nearby ones [Hoyle, 1960; Sandage, 1961; Ellis, 1971a]. For the Einstein-de Sitter universe, the minimum angular diameter is at $z_* = 1.25$; in low density universes, it occurs at higher redshifts.

The number of objects seen in a solid angle $d\Omega$ for a distance increment du (characterized by increments dz , dm in the observable variables z and m) is given by

$$(27) \quad dN = W(u) \rho(t_e) S^3(t_e) f(u) du d\Omega$$

where the detection probability or ‘selection function’ is $W(u)$ [Dodelson, 2003, p. 263] and $\rho(t_e)$ is the number density of objects at the time of emission (spatial homogeneity is expressed in the fact that this is independent of the spatial coordinates). The observed total number N of objects measured in a survey is given by integrating from the observer to the survey limit: in terms of the radial coordinate r_e of the source (which can be related to redshifts or magnitudes), $N = \int_{r_0}^{r_e} dN$. If the number of objects is conserved (e.g. observing galaxies in an epoch when they are neither created nor destroyed), $\rho(t_e) = \rho(t_0)(1 + z)^3$ and we find from (27) that in the idealized case when W is independent of distance (a reasonable assumption for relatively nearby objects),

$$(28) \quad N = W \rho(t_0) d\Omega \int_{r_0}^{r_e} f(u) du.$$

The simple integral has to be separately done for the cases $k = +1, 0, -1$ [Sandage, 1961].

The above equations enable one to determine observational relations between observable variables, for example (m, z) , (α, z) or (N, m) relations for objects with known intrinsic properties (known size or luminosity, for example), and so

¹³Or, more elegantly, from the geodesic deviation equation (see [Ellis and van Elst, 1999b]).

to observationally determine q_0 . These relations have to be modified if there is absorption by an intergalactic medium, gravitational lensing, or anisotropic emission of radiation; and detailed comparisons with observations have to take into account the spectrum of the source as well as source detection and identification probabilities [Harwit, 1984]. Here we encounter *the contrast between image and reality*: there can be many objects out there that we either do not detect, or do not recognize for what they are [Disney, 1976]. An “observational map” relating source properties to the nature of their images gives a useful view of how this occurs [Ellis *et al.*, 1984].

One important feature here is that *a specific object will look completely different at different wavelengths* (optical, radio, X-ray for example); indeed it may be detectable at one wavelength and not at another. This shows very clearly how our images of reality are dependent on the detectors we use. To get a full picture of what is out there, we need to use multiple modes of investigation — imaging at all wavelengths together with intensity, spectral, and polarization measurements [Harwit, 1984], as well as watching for time variations. A second important feature is observational selection effects such as the *Malmquist bias* — if we have a population of objects with different luminosities, at large distances we will only see the more luminous objects (the fainter ones will not be detected); hence *the average luminosity will appear to increase with distance*, but this is just an observational effect rather than the real state of affairs. Using different detection thresholds controls this effect to some degree.

2.3.4 Number Counts and the visible matter density

Number counts of galaxies as a function of redshift or luminosity show approximate spatial homogeneity of the universe [Hubble, 1936]. However counts of radio sources and quasi-stellar objects (qso’s) show that the universe has not been in a steady state as proposed by Bondi, Gold, and Hoyle [Bondi, 1960]. Indeed *number counts are only compatible with a RW geometry if there has been evolution of source numbers and/or luminosities* [Sciama, 1971].

Number counts also give estimates of the density of visible (luminous) matter in the universe: $\Omega_{vm0} \simeq 0.015$. This is very low relative to the critical density ($\Omega_0 = 1$) and is also considerably less than the amount of baryons determined by nucleosynthesis studies ($\Omega_{bar0} \simeq 0.044$, see Sec. 2.2.2). Thus *much of the baryonic matter in the universe is in some hidden (non-luminous) form* [Bothun, 1998, pp. 223-272], e.g. burnt out stars [Hogan, 1999].

2.3.5 Apparent Luminosities and sizes: Dark Energy

Apparent sizes or luminosities as a function of redshift can be used to determine the deceleration parameter q_0 (Sec. 2.1) if the intrinsic source sizes or luminosities are known. The problem is that until recently there were not known enough galaxies or other objects of standard size or luminosity to use to determine q_0 , and scatter in their properties leads to biasing of observations by the Malmquist

effect. However this dramatically changed with recent observations of the decay curves of the luminosity of supernovae in distant galaxies. It turns out that the peak luminosity of type Ia supernovae is closely correlated with their light curve decay time, for the first time giving reliable ‘standard candles’ for galaxies at large distances [Perlmutter *et al.*, 1998]. The conclusion from these observations is that, rather than slowing down as expected, *the rate of expansion of the universe is speeding up at a rate corresponding to a cosmological constant with $\Omega_{\Lambda 0} = 0.7$* . This evidence is concordant with that from CBR observations and number counts [Dodelson, 2003; Silk, 2005].

The nature of the field or matter causing this acceleration is unclear. Its equation of state $w := pc^2/\mu$, is unknown, and many physical and unphysical proposals are being made in this regard. From (7), it has to violate the strong energy condition (8) and so must have a large negative pressure. It could indeed be due to a cosmological constant ($w = -1$), which would have dominated the expansion of the universe since a redshift $z \simeq 0.33$, and would have been negligible earlier (and is also negligible on small scales — it does not affect local astrophysics). However it could also be some other form of matter or field with effective negative pressure so that $w < -1/3$, as can happen in the case of a scalar field (see eqn.(33) below). In that case it is called ‘quintessence’. There are many speculations as to what this might be, but there is no clarity on the matter. One should note here that alternative explanations of the observations are possible, for they can be exactly reproduced by a spherically symmetric inhomogeneous universe model where we are near the centre [Mustapha *et al.*, 1998], or could at least partly be due to the back-reaction of inhomogeneities on the cosmic expansion [Ellis and Buchert, 2006] or the effect of inhomogeneities on the effective area distance [Kantowski *et al.*, 1995; Kantowski, 1998]. These alternatives are being investigated, but the most probable cause remains some unknown kind of matter or field with effective negative energies.

In summary, the standard gravitational equations together with the supernovae observations imply *presence of a cosmological constant or some equivalent form of ‘dark energy’ with a large effective negative energy density μ_{grav} (due to negative pressure) dominating the present expansion of the universe; its physical nature is unknown*. There is no known physics reason why this force should exist at this level where it is just detectable — quantum field theory relates the cosmological constant to the zero-point energy of the vacuum, and suggests it should be enormously larger than observed [Weinberg, 1989; 2000a; 2000b; Rugh and Zinkernagel, 2002; Zinkernagel, 2002; Susskind, 2005]. It is a major mystery why it exists at the small (just detectable) level that observations indicate [Seife, 2003]. A key aspect of present day cosmology is trying on the one hand to observationally determine the effective equation of state of this ‘dark energy’ (running the field equations backwards to obtain $w(z)$ from the observations [Saini *et al.*, 2000], and in particular determining whether w is constant or varying over time), and on the other attempting to give a plausible theoretical explanation for its physical origin.

2.3.6 Matter Distribution and Motion: Dark Matter

Detailed studies have been made of the distribution of galaxies and their motions. They occur in clusters, in turn making up superclusters imbedded in vast walls surrounding relatively sparsely populated intergalactic voids. The galaxy *luminosity function* characterizes the numbers of galaxies occurring within each luminosity class; the *covariance function* characterizes their spatial clustering [Peebles, 1993b; Dodelson, 2003]. Large scale motions occur for galaxies in clusters, and for the clusters themselves. It is easy to conceive of matter that is hard to detect (for example, small rocks distributed through space); studies of galactic rotation curves and of motions of galaxies in clusters [Bothun, 1998, pp. 139-181] imply *the existence of huge amounts of unseen dark matter, dominating the dynamics of the Universe*: its density is $\Omega_{dm0} \simeq 0.3$, much greater than both visible matter ($\Omega_{vm0} = 0.015$) and baryons ($\Omega_{bar0} = 0.044$), but significantly less than the critical density $\Omega_0 = 1$. Thus *the dark matter is non-baryonic*, meaning it has some kind of exotic nature rather than being the protons and neutrons that are the substance of ordinary matter [Seife, 2003]. In contrast to the ‘dark energy’ discussed in the previous section, dark matter is dynamically effective on astrophysical scales as well as on cosmological scales. Many attempts have been made to identify its nature, for example it might be axions, supersymmetric partners of known particles, quark nuggets, or massive neutrinos [Gribbin and Rees, 1991; Perkins, 2005], but what it is composed of is still unknown. Laboratory searches are under way to try to detect this matter, so far without success. A key question is whether its apparent existence is due to our using the wrong theory of gravity on these scales. This is under investigation, with various proposals for modified forms of the gravitational equations that might explain the observations without the presence of large quantities of dark matter. This is a theoretical possibility, but the current consensus is that this dark matter does indeed exist.

An important distinction is whether dark matter consists of

- (i) weakly interacting massive particles that cooled down quickly, thereafter forming *cold dark matter* (‘CDM’) moving slowly at the time of structure formation (and resulting in a bottom-up process with large scale structure forming from smaller scale structures), or
- (ii) particles that have a low mass and cooled slowly, therefore for a long time forming *hot dark matter*, moving very fast at the time of structure formation (and resulting in a top-down galaxy formation scenario).

Structure formation studies currently support the CDM hypothesis, with hierarchical formation of gravitationally bound objects taking place in a complex bottom up process involving interactions of CDM, baryons, and radiation, with dwarf galaxies forming initially [Silk, 2005; Mouri and Taniguchi, 2005] and then aggregating to form larger structures. These studies are based on massive numerical simulations, with initial conditions provided by the inflationary scenario discussed below, see

Sec. 2.6. Unlike ‘dark energy’, CDM has an ordinary baryonic equation of state (it is a perfect fluid (4) with $p_{cdm} = 0 \Leftrightarrow w_{cdm} = 0$).

Another way of detecting dark matter in clusters is by its gravitational lensing effects [Schneider *et al.*, 1992]. The bending of light by massive objects was one of the classic predictions of General Relativity theory. Rich clusters of galaxies or galaxy cores can cause strong lensing of more distant objects, where multiple images of distance galaxies and qso’s occur, sometimes forming rings or arcs; and weaker lensing by closer masses results in characteristic patterns of distortions of images of distant objects. Analysis of multiple images can be used to reconstruct the lensing mass distributions, and statistical analysis of weak lensing patterns of image distortions are now giving us detailed information on the matter distribution in distant galaxies and clusters. These studies show that to get enough lenses in an almost flat cosmology ($\Omega_0 \simeq 1$) requires the presence of a cosmological constant — there cannot be a critical density of dark matter present [Dodelson, 2003; Silk, 2005].

A key feature of present-day cosmology is attempts to identify the nature of this dark matter, and if possible to detect it in a laboratory situation. While observations favour the CDM scenario, some residual problems as regards the emergence of fine-scale structure still need resolution [Silk, 2005].

2.3.7 The CBR Power spectrum

The CBR angular anisotropies are characterized by an angular power spectrum showing the amount of power in perturbations at each physical scale on the LSS [Bennet *et al.*, 2003; Seife, 2003; Dodelson, 2003]. In the time from the the end of inflation to the LSS, modes of different wavelengths can complete a different number of oscillations. This translates the characteristic periods of the waves into characteristic lengths on the LSS, leading to a series of maxima (‘acoustic peaks’) and minima in the inhomogeneities on the LSS and consequently in the CBR angular anisotropy power spectrum [Hu and Sugiyama, 1995b; Hu and Sugiyama, 1995a; Peacock, 1999; Perkins, 2005]. These inhomogeneities then form the seeds for structure formation and so are related to the power spectrum of physical scales for structures that form later. They are characterised by a (3-dimensional) spatial power spectrum on the LSS; because we receive the observed CBR radiation from the 2-sphere $S_{2:LSS}$ where our past light cone intersects the LSS, this is seen by us as a 2-dimensional power spectrum of anisotropies on the sky (characterised by the unit sphere S_2 of all direction vectors e_a : $e^a e_a = 1$, $e^a u_a = 0$).

The apparent angular size of the largest CBR peak (about 1°) allows estimates of the area distance to the LSS and hence of the density of matter in the universe for various values of the cosmological constant, and determines the major cosmological parameters [Spergel *et al.*, 2003]:

“By combining WMAP data with other astronomical data sets, we constrain the geometry of the universe: $\Omega_{tot} = 1.02 \pm 0.02$, the equation of state of the dark energy, $w < -0.78$ (95% confidence limit), and the

energy density in neutrinos, $\Omega_\nu h^2 < 0.0076$ (95% confidence limit). For 3 degenerate neutrino species, this limit implies that their mass is less than 0.23 eV (95% confidence limit). The WMAP detection of early reionization rules out warm dark matter.”

There is however a problem here: while the agreement of theory and observations for all small angular scales is remarkable, there is a divergence at the largest angular scales: the observations show less power than expected. Specifically, the quadrupole and octopole are much lower than theory predicts. Also the axes of the quadrupole and octopole are very precisely aligned with each other, and there are other angular anomalies [Starkman and Schwarz, 2005]. These effects might be due to (i) observational contamination by the galaxy (which gets in the way of our view of the LSS), (ii) a contingent (‘chance’) event (it represents ‘cosmic variance’, discussed below, see Sec. 3), (iii) our living in a well-proportioned ‘small universe’ which is spatially closed so that there is a maximum size to possible fluctuations [Weeks *et al.*, 2003], or (iv) some unexpected new physical effect or deeper problem with our understanding of the early universe. The jury is out as to which the case is; this could turn out to be a crisis for the CBR analysis, but on the other hand one can always just resort to saying it is a statistical fluke (the underlying problem here being the uniqueness of the universe, as discussed in Sec. 3).

There are similar expected peaks in the polarization spectrum of this radiation, and polarization maps should have a mode associated with gravitational waves predicted by inflation to exist in the very early universe (Sec. 2.6); detection of such modes will be a crucial test of inflation [Dodelson, 2003; Sievers *et al.*, 2005b]. Studies of polarization indicate that reionisation of the universe took place as early as a redshift of 17, contrary to what is deduced from qso studies. More detailed studies of anisotropies involve the Sunyaev-Zel’dovich effect (changes in the observed temperature due to scattering by hot matter in galaxy clusters) and gravitational lensing.

There is a huge amount of information in the CBR maps, and their more accurate measurement and interpretation is a central feature of current cosmology [Steinhardt, 1995; Peacock, 1999; Dodelson, 2003; Perkins, 2005].

2.4 Causal and visual horizons

A fundamental feature affecting the formation of structure and our observational situation is the limits arising because causal influences cannot propagate at speeds greater than the speed of light. Thus the region that can causally influence us is bounded by our past null cone. Combined with the finite age of the universe, this leads to the existence of particle horizons limiting the part of the universe with which we can have had causal connection.¹⁴

A *particle horizon* is by definition comprised by the limiting worldlines of the furthest matter that ever intersects our past null cone [Rindler, 1956; 2001]. This

¹⁴There are also *event horizons* and *apparent horizons* in some cosmological models [Rindler, 1956; Tipler *et al.*, 1980] and [Rindler, 2001, pp. 376-383].

is the limit of matter that can have had any kind of causal contact with since the start of the universe, characterized by the comoving radial coordinate value

$$(29) \quad u_{ph} = \int_0^{t_0} \frac{dt}{S(t)}.$$

The present physical distance to the matter comprising the horizon is

$$(30) \quad d_{ph} = S(t_0)u_{ph}.$$

The key question is whether the integral (29) converges or diverges as we go to the limit of the initial singularity where $S \rightarrow 0$. *Horizons will exist in standard FL cosmologies for all ordinary matter and radiation*, for u_{ph} will be finite in those cases; for example in the Einstein-de Sitter universe (see Sec. 2.1.1), $u_{ph} = 3t_0^{1/3}$, $d_{ph} = 3t_0$. We will then have seen only a fraction of what exists, unless we live in a universe with spatially compact sections so small that light has indeed had time to traverse the whole universe since its start; this will not be true for universes with the standard simply-connected topology. Penrose's powerful use of conformal methods (see [Hawking and Ellis, 1973; Tipler *et al.*, 1980]) gives a very clear geometrical picture of the nature of horizons [Ellis and Williams, 2000]. They may not exist in non-FL universes, for example Bianchi (anisotropic) models [Misner, 1969]. In universes with closed spatial sections, a supplementary question arises: Is the scale of closure smaller than the horizon scale? There may be a finite time when causal connectivity is attained, and particle horizons cease to exist. In standard $k = +1$ FL models, this occurs just as the universe reaches the final singularity; if however there is a positive cosmological constant or other effective positive energy density field, it will occur earlier. The horizon always grows, because (29) shows that u_{ph} is a monotonically increasing function of t_0 . Despite many contrary statements in the literature, *it is not possible that matter leave the horizon once it has entered*. In a (perturbed) FL model, once causal contact has taken place, it remains until the end of the universe.

The importance of horizons is two-fold: they underlie causal limitations relevant in the origin of structure and uniformity [Misner, 1969; Guth, 1981], and they represent absolute limits on what is testable in the universe [Ellis, 1975; 1980].

2.4.1 Causal limitations

As to causal limitations, horizons are important in regard both to the smoothness of the universe on large scales, and the lumpiness of the universe on small scales. The issue of smoothness is encapsulated in the *horizon problem*: if we measure the temperature of the CBR arriving here from opposite directions in the sky in a standard FL model, it came from regions of the surface of last scattering that can have had no causal contact of any kind with each other since the start of the universe. In a radiation-dominated early universe with scale factor (19), the size of the particle horizon at the time of last scattering appears as an angular scale of about 1° in the sky today, and corresponds to a comoving physical length of about

400,000 light years when evaluated today. Why then are conditions so similar in these widely separated regions? [Misner, 1968; Guth, 1981; Blau and Guth, 1987; Kolb and Turner, 1990]. Note that this question is of a philosophical rather than physical nature, i.e. there is no contradiction here with any experiment, but rather an unease with an apparent fine tuning in initial conditions. This problem is claimed to be solved by the inflationary universe scenario mentioned below, see Sec. 2.6.

Associated with the existence of horizons is the prediction that physical fields in different regions in the universe should be uncorrelated after symmetry breaking takes place, because they cannot have interacted causally. Consequently, if grand unified theories are correct, topological defects such as monopoles and cosmic strings may be expected as relics of the expansion of the very early universe [Kolb and Turner, 1990]. In a standard cosmology, far too many monopoles are predicted. This is also solved by inflation.

As to the lumpiness, the issue here is that if we believe there was a state of the universe that was very smooth — as indicated at the time of decoupling, by the low degree of anisotropy of the CBR, and represented by the RW geometry of the FL models — then there are limits to the sizes of structures that can have grown since then by causal physical processes, and to the relative velocities of motion that can have been caused by gravitational attraction in the available time (for example, the peculiar motion of our own galaxy relative to the CBR rest frame caused by the huge overdensity called the ‘Great Attractor’). If there are larger scale structures or higher velocities, these must have been imprinted in the perturbations at the time of last scattering, for they cannot have been generated in a causal way since that time. They are set into the initial conditions, rather than having arisen by physical causation from a more uniform situation. This is a key factor in the theory of growth of perturbations in the early universe where the expansion damps their growth. The quantity determining the relevant physical scales for local causal influences in an expanding universe is the *comoving Hubble radius* $\lambda_H := (SH)^{-1}$; the way perturbations of wavelength λ develop depends on whether $\lambda > \lambda_H$ or $\lambda < \lambda_H$ [Dodelson, 2003, pp. 146-150].

Actually the domain of causal influence is even more tightly constricted than indicated by the past light cone: the limits coming from the horizon size are limits on what can be influenced by particles and forces acting at the speed of light. However only freely travelling photons, massless neutrinos, and gravitons can move at that speed; and such particles coming from cosmological distances have very little influence on our galaxy or the solar system (indeed we need very delicate experiments to detect them). Any massive particles, or massless particles that are interacting with matter, will travel slower (for example before decoupling, light has a very small mean free path and information will travel only by sound waves and diffusion in the tightly coupled matter-radiation fluid). The characteristics for pressure-free scalar and vector perturbations are timelike curves, moving at zero velocity relative to the matter; while density perturbations with pressure can move at the speed of sound, only tensor perturbations can travel at the speed

of light. Thus the true domain that influences us significantly is much less than indicated by the particle horizon. It is the small region round our past world line characterised after decoupling by the comoving scale from which matter coalesced into our galaxy: a present distance of about 1 to 1.95 Mpc,¹⁵ corresponding to an observed angle of about 0.64 arcminutes on the LSS. Before decoupling it would have been limited by the sound horizon [Dodelson, 2003, p. 257] rather than the particle horizon.

2.4.2 Observational limitations

Clearly we cannot obtain any observational data on what is happening beyond the particle horizon; indeed we cannot even see that far because the universe was opaque before decoupling. *Our view of the universe is limited by the visual horizon, comprised of the worldlines of furthest matter we can observe — namely, the matter that emitted the CBR at the time of last scattering* [Ellis and Stoeger, 1988]. This occurred at the time of decoupling $t = t_{dec}$ ($z_{dec} \simeq 1100$), and so the visual horizon is characterized by $r = u_{vh}$ where

$$(31) \quad u_{vh} = \int_{t_{dec}}^{t_0} \frac{dt}{S(t)} < u_{ph}.$$

Indeed the LSS delineates our visual horizon in two ways: we are unable to see to *earlier times* than its occurrence (because the early universe was opaque for $t < t_{dec}$), and we are unable to detect matter at *larger distances* than that we see on the LSS (we cannot receive radiation from matter at co-moving coordinate values $r > u_{vh}$). The picture we obtain of the LSS by measuring the CBR from satellites such as COBE and WMAP is just a view of the matter comprising the visual horizon, viewed by us at the time in the far distant past when it decoupled from radiation. The position of the visual horizon is determined by the geometry since decoupling. Visual horizons do indeed exist, unless we live in a small universe, spatially closed with the closure scale so small that we can have seen right around the universe since decoupling. This is a possibility that will be discussed below (Sec. 4.3.1). There is no change in these visual horizons if there was an early inflationary period, for inflation does not affect the expansion or null geodesics during this later period. The major consequence of the existence of visual horizons is that many present-day speculations about the super-horizon structure of the universe — e.g. the chaotic inflationary theory (Sec. 2.6) — are not observationally testable, because one can obtain no definite information whatever about what lies beyond the visual horizon [Ellis, 1975; 1980]. This is one of the major limits to be taken into account in our attempts to test the veracity of cosmological models (Sec. 4.3).

¹⁵Dodelson [Dodelson, 2003], p. 283; W Stoeger, private communication.

2.5 Theoretical Developments

The cosmological application of Einstein's Theory of Gravitation has also progressed greatly in past decades, as regards exact solutions and generic properties of the field equations; as regards approximate solutions; and in terms of understanding the relationship between them.

2.5.1 Exact solutions and generic properties

Theory initially predicted there must have been a start to the universe, but it was not clear for a long time if this was simply due to the very special exactly isotropic and spatially homogeneous geometry of the standard Friedmann-Lemaître models. It was possible that more realistic models with rotation and acceleration might show the prediction was a mathematical artefact resulting from the idealized models used. The singularity theorems developed by Penrose and Hawking [Hawking and Ellis, 1973; Tipler *et al.*, 1980; Earman, 1999] showed this was not the case: even for realistic geometries, classical gravitational theory predicts a beginning to the universe at a space-time singularity, provided the usual energy conditions were satisfied. This study has led inter alia to a greatly increased understanding of causality and topology of generic universe models [Tipler *et al.*, 1980], including the fact that singularities may have a quite different nature than those in the Robertson-Walker models, for example being anisotropic [Tipler *et al.*, 1980] or of a non-scalar nature [Ellis and King, 1974].

Various classes of exact cosmological solutions are known (Kantowski-Sachs and Bianchi spatially homogeneous but anisotropic models, Tolman-Bondi spherically symmetric inhomogeneous models, and 'Swiss-Cheese' non-analytic models) enabling understanding of dynamical and observational behaviour of more general classes of models than just the FL models [Ellis and van Elst, 1999a]. Dynamical systems studies [Wainwright and Ellis, 1996; Uggla *et al.*, 2003] relate the behaviour of whole classes of anisotropic models in suitable state spaces, enabling identification of generic patterns of behaviour (fixed points, saddle points, attractors, etc.) and hence the relationship between dynamics of higher symmetry and lower symmetry universes. These studies help understanding to what degree the FL models are generic within the families of possible cosmological models, and which models might give observations similar to those in the FL models. In particular they are relevant in considering the possible geometry of the universe at very early or very late times.

2.5.2 Perturbation theory, the gauge issue, and back reaction

Sophisticated perturbation theory has been developed to underlie the theory of structure formation in the expanding universe, examining *the dynamics of perturbed FL models*. The fluid flow in these models can have shear, vorticity, and acceleration, and the Weyl tensor C_{ijkl} (see (1)) is not zero, so that density variations, tidal forces, peculiar velocities, and gravitational waves can be present. De-

tailed studies use the kinetic theory approximation for matter (electrons, protons, dark matter) and radiation (photons, neutrinos), with their dynamics described by the Boltzmann equation [Dodelson, 2003, Ch. 4]; [Uffink, 2006], interacting with space-time inhomogeneities characterised by a perturbed FL metric. A key issue here is the *gauge problem* — how to choose the background model in the perturbed spacetime [Ellis and Stoeger, 1987]. If this is not properly handled then one may attain apparent perturbation solutions that are pure gauge (they are mathematical rather than physical), so that one can alter the apparent growth rate simply by changing coordinates. The key to handling this is either to keep careful track at all stages of remaining gauge freedom and possible changes of gauge, or (preferably, in my view) to use gauge invariant variables (see [Bardeen, 1980; Ellis and Bruni, 1989; Challinor and Lasenby, 1998]).

Most of the literature on perturbation theory deals with the linear case, but some studies tackle the non-linear regime (e.g. [Langlois and Vernizzi, 2005]), and some consider questions such as the origin of magnetic fields and the causes of galactic rotation. A key problem here is properly relating relativistic analyses of astrophysical dynamics to the Newtonian approaches most often used by astrophysicists (e.g. [Bothun, 1998], pp. 183-222); this is not straightforward.¹⁶ A further unresolved issue is the nature of gravitational entropy [Penrose, 1989b; Ellis, 2002; Penrose, 2004]. Many statements about the nature of entropy in physics textbooks are wrong when gravity is dominant, leading to the spontaneous formation of structures such as stars and galaxies. There is as yet no agreed definition of gravitational entropy that is generally applicable; until there is, cosmological arguments relying on entropy concepts are ill-founded.

The existence of inhomogeneities in the universe raises the issue of fitting and back-reaction. To what degree does the nature of the exactly smooth FL models reflect the geometrical and dynamical nature of more realistic ‘lumpy’ universe models? [Ellis and Stoeger, 1987]. Inhomogeneities lead to extra terms appearing in the evolution equations for the idealized background models, representing the back-reaction of the perturbations on their dynamics [Ellis, 1984]. These could possibly be dynamically significant [Ellis and Buchert, 2006], but this is a matter of dispute.

2.6 Inflation

Particle physics processes dominated the very early eras, when exotic processes took place such as the condensation of a quark-gluon plasma to produce baryons. Quantum field theory effects were significant then, and this leads to an important possibility: scalar fields producing repulsive gravitational effects could have

¹⁶Some exact General Relativity results, which must necessarily apply in the Newtonian limit of General Relativity, have no Newtonian analogue; an example is the shear-free theorem applying to pressure-free matter [Ellis, 1967]. The underlying issue is that there are 10 field equations to be satisfied in General Relativity, with 20 integrability conditions (the Bianchi identities), but only one field equation to be satisfied in Newtonian theory (Poisson’s equation) together with 4 conservation equations.

dominated the dynamics of the universe at those times. This leads to the theory of the inflationary universe, proposed by Alan Guth [1981; 1997]: if $\mu_{grav} = \mu + 3p/c^2 < 0$, which can happen if a scalar field dominates the dynamics of the early universe, an extremely short period of accelerating expansion will precede the hot big bang era [Blau and Guth, 1987]. This produces a very cold and smooth vacuum-dominated state, and ends in ‘reheating’: conversion of the scalar field to radiation, initiating the hot big bang epoch. This inflationary process is claimed to explain the puzzles mentioned above (Sec. 2.4.1): why the universe is so special (with spatially homogeneous and isotropic geometry and a very uniform distribution of matter), and also why the space sections are so close to being flat at present (we still do not know the sign of the spatial curvature), which requires very fine tuning of initial conditions at very early times. Inflationary expansion explains these features because particle horizons in inflationary FL models will be much larger than in the standard models with ordinary matter, allowing causal connection of matter on scales larger than the visual horizon, and inflation also will sweep topological defects outside the visible domain.

In more detail: in the case of a single scalar field ϕ with spacelike surfaces of constant density, on choosing u^a orthogonal to these surfaces, the stress tensor has a perfect fluid form with

$$(32) \quad \mu = \frac{1}{2}\dot{\phi}^2 + V(\phi), \quad p/c^2 = \frac{1}{2}\dot{\phi}^2 - V(\phi),$$

and so

$$(33) \quad \mu + 3p/c^2 = 2\dot{\phi}^2 - 2V(\phi).$$

The slow-rolling case is $\dot{\phi}^2 \ll V(\phi)$, leading to $\mu + p/c^2 = 2\dot{\phi}^2 \simeq 0 \Rightarrow \mu + 3p/c^2 \simeq -2\mu < 0$. This then enables a resolution of the horizon problem in inflationary FL models: if sufficient inflation took place in the early universe, then all the regions from which we receive CBR were causally connected; indeed if the universe began in an inflationary state, or was inflationary with compact spatial sections, there may be no causal horizons at all. The inflationary models also cause initial perturbations to die away, including velocity perturbations, hence explaining the observed smoothness of the universe on large scales. This process is expected to create a universe with spatially very flat sections at late times:

$$(34) \quad \Omega_0 = \Omega_{dm0} + \Omega_{\Lambda 0} \simeq 1 \Leftrightarrow \Omega_k \simeq 0.$$

This theory led to a major bonus: a proposal that initial tiny quantum fluctuations were expanded to such a large scale by inflation that they provided seeds initiating growth by gravitational attraction of large scale structures such as clusters of galaxies. This theory makes clear observational predictions for the spectrum of CBR anisotropies, which have since been spectacularly verified by observations from balloons and satellites, such as WMAP [Spergel *et al.*, 2003]. Thus inflation has provided us with our first coherent theory of structure formation. Inhomogeneities started as quantum fluctuations in the inflationary

epoch which are then amplified in physical scale by the inflationary expansion but remain constant in amplitude when larger than the contemporary Hubble scale, leading to Gaussian scale-free perturbations at the start of the HBB era. Starting from these fluctuations, Cold Dark Matter ('CDM') creates potential wells for baryons to fall into, but the radiation (tightly coupled to the electrons and baryons) resists collapse. Gravity wins if the wavelength λ is greater than the *Jean's length* λ_J (which is proportional to the speed of sound [Rees, 1995; Ellis and van Elst, 1999a]). There are acoustic oscillations (sound waves) when $\lambda < \lambda_J$; these oscillations ceased at decoupling, which led to a dramatic decrease in λ_J and the growth of structure by gravitational instability in a 'bottom up' way (Sec. 2.3.6).¹⁷

A popular version of inflation is *chaotic inflation* [Linde, 1990; Guth, 2001; Susskind, 2005] where inflation ends at different times in different places, so that one ends up with numerous 'pocket universe' (expanding universe domains like the one we see around us, or perhaps very different) all imbedded in a still-inflating universe region and starting at different times, the whole forming a fractal-like structure. It is argued this is an inevitable consequence of the nature of plausible scalar field potentials.

Inflation is not an inevitable conclusion, for there are some alternatives proposed [Hollands and Wald, 2002; Khoury *et al.*, 2001], and the WMAP results can be reproduced by any scenario where Gaussian scale-free perturbations of suitable amplitude occur at the start of the Hot Big Bang era. However inflation is regarded by most cosmologists as the best proposal available for the era prior to the Hot Big Bang epoch, leading to the presence of such perturbations. Nevertheless one should note it is a generic proposal for what happened, rather than a specific physical theory. While a great many possibilities have been proposed (it could for example be an effective field due to higher-order gravity effects, or it could involve multiple scalar fields), at the present time the identity of the proposed inflationary field ('the inflaton') has not been established or linked to any known particle or field. The hoped-for link between early universe dynamics and particle physics is potential rather than real [Earman and Mosterin, 1999]. Detailed studies of the CBR anisotropies and structure formation in conjunction with the observations hope to distinguish between the various possibilities, for example testing whether the spectral index n takes the scale-free value: $n = 1$, or whether rather there is a tilted power spectrum ($n \neq 1$). A unique spectrum of gravitational waves will also be produced at very early times in an inflationary universe, and detection of these waves either directly by proposed gravitational wave detectors or indirectly by measuring the associated curl mode in the CBR polarization will be an important test of inflation, for example determining the ratio r of scalar to tensor perturbations in the early universe [Dodelson, 2003].

¹⁷This is a highly simplified account; for more detailed versions, see e.g. [Dodelson, 2003; Silk, 2005].

2.7 The very early universe

Quantum gravity processes are presumed to have dominated the very earliest times, preceding inflation. There are many theories of the quantum origin of the universe, but none has attained dominance. The problem is that we do not have a good theory of quantum gravity [Rovelli, 2006], so all these attempts are essentially different proposals for extrapolating known physics into the unknown. A key issue is whether quantum effects can remove the initial singularity and make possible universes without a beginning. Preliminary results suggest this may be so [Bojowald, 2001; Rovelli, 2004; Mulryne *et al.*, 2005].

2.7.1 Is there a quantum gravity epoch?

A preliminary issue is, can there be a non-singular start to the inflationary era, thus avoiding the need to contemplate a preceding quantum gravity epoch? In the inflationary epoch the existence of an effective scalar field leads to a violation of the strong energy condition (8), therefore at first sight it seems that a bounce may be possible preceding the start of the expanding inflationary era and avoiding the inevitability of a quantum gravity epoch.

However a series of theorems suggest that inflationary models cannot bounce: they are stated to be future infinite but not past infinite [Guth, 2001]. This is an important issue, so it is worth looking at it further. There are two major requirements to get a bounce. The Friedmann equation (9) relates the scale factor $S(t)$, curvature constant k , and the effective total energy density $\mu(t)$, which is *defined* by this equation whatever dynamics may be involved (multiple scalar fields, higher order gravity, higher dimensional theories leading to effective 4-dimensional theories, etc.).¹⁸ The Raychaudhuri equation (7) includes the effective total pressure $p(t)$, which again is *defined* by this equation. In this section, a cosmological constant Λ is represented as perfect fluid with $\mu_\Lambda + p_\Lambda/c^2 = 0$. To get a bounce, first one needs the curve $S(t)$ of the scale factor as a function of time to bend up: that is,

$$(35) \quad \frac{\ddot{S}}{S} \geq 0 \Leftrightarrow \mu + 3p/c^2 < 0,$$

which is just a violation of the strong energy condition (8). This is the case if $\mu + p/c^2 = 0$ (a vacuum); and indeed by eqn.(33) it is possible for example for any slow-rolling scalar field. Second, one also needs a time when the scale factor is a minimum. Thus there must be a time t_* such that $\dot{S}(t_*) = 0$. From the Friedmann equation (9),

$$(36) \quad \dot{S}^2(t_*) = 0 \Leftrightarrow \frac{\kappa\mu(t_*)}{3} = \frac{k}{S^2(t_*)}.$$

¹⁸See [Copeland *et al.*, 2005] for the ways various quantum gravity theories result in modified Friedmann equations.

With $k \leq 0$ this is possible only if $\mu(t_*) < 0$. Even with a scalar field (see eqn.(32)) this can only be achieved by having negative potential energies, which appears to be an unphysical requirement. With $k = +1$ this is possible with $\mu(t_*) > 0$ [Robertson, 1933], which is compatible with ordinary matter.

Thus if you want a bounce in an inflationary universe, it is sensible to look to $k = +1$ inflationary models, which indeed will turn around if a vacuum domain occurs for long enough (curvature will eventually always win over a vacuum as we go back into the past [Ellis *et al.*, 2002b; Ellis *et al.*, 2002a]). The theorems mentioned above do *not* include this case (see [Guth, 2001]); they only consider inflationary universes with $k = 0$ and $k = -1$. And one should note here that although the scale-free $k = 0$ exponential case clearly is the model underlying the way many people approach the problem, it is highly exceptional — it is of zero measure within the space of all inflationary FL models.

Explicit non-singular models can be constructed, the simplest being the de Sitter universe in the $k = +1$ frame (Sec. 2.1.1), which is an exact eternal solution that bounces at a minimum radius S_0 . This model has the problem that it does not exit inflation (it corresponds to an exactly constant potential), but variants exist where exit is possible; there are also viable non-singular models that start off asymptotic to the Einstein Static universe in the distant past and avoid the need for a quantum gravity epoch [Ellis and Maartens, 2004]. These models start off in a very special state, precisely because they are asymptotic to the Einstein static universe in the distant past. This is a possible situation. It seems likely that the options for the start of inflation are (i) avoiding the quantum gravity era, but at the cost of having special ('fine tuned') initial conditions, or (ii) having a quantum gravity epoch preceding the inflationary era. Thus a key issue is whether the start of the universe was very special or generic.

2.7.2 *Quantum gravity effects: The origin of the universe*

Contemporary efforts to explain the beginning of the universe, and the particular initial conditions that have shaped its evolution, usually adopt some approach or other to applying quantum theory to the creation of the universe [Lemaître, 1931]. Many innovative attempts have been made here; as this article focuses on General Relativity and its application to cosmology, and it would be impossible to do justice to the various approaches to quantum cosmology [Rovelli, 2006] without a very much longer article. I will just make a few comments on these approaches.

The attempt to develop a fully adequate quantum gravity approach to cosmology is of course hampered by the lack of a fully adequate theory of quantum gravity, as well as by the problems at the foundation of quantum theory (the measurement problem, collapse of the wave function, etc., see [Isham, 1997; Dickson, 2006; Landsman, 2006]) which can be ignored in many laboratory situations but have to be faced in the cosmological context [Perez *et al.*, 2005]. The various attempts at quantum cosmology each develop in depth some specific aspect of quantum theory that may be expected to emerge from a successful theory of quantum grav-

ity applied to the universe as a whole, being for example based on either (i) the Wheeler-DeWitt equation and the idea of the wave function of the universe, or (ii) on some version of embedding in higher dimensional space time (inspired by string theory), or (iii) an appropriate application of loop quantum gravity. In effect they attempt either

- (a) to give a true theory of creation *ex nihilo* [Vilenkin, 1982]; such efforts however cannot truly “solve” the issue of creation, for they rely on some structures or other (e.g. the elaborate framework of quantum field theory and much of the standard model of particle physics) pre-existing the origin of the universe, and hence themselves requiring explanation; or
- (b) to describe a self-sustaining or self-referential universe which by-passes the issue of creation, either by
 - (b1) originating from an eternally pre-existing state, via the recurring idea of a Phoenix universe [Dicke and Peebles, 1979] (as in Veneziano’s ‘pre-big bang theory’ based on analogues of the dualities of string theory, or self-repeating universes such as the chaotic inflationary models of Linde); creation from fluctuations in some quite different pre-existing structure (e.g. emergence from de Sitter space time; or the ‘ekpyrotic universe’ initiated by a collision between pre-existing ‘branes’ in a higher dimensional spacetime); or emerging from an eternal static initial state; or
 - (b2) starting from a state with different properties of time than usual (or with an emergent notion of time): as in the Hartle–Hawking no-boundary proposal [Hawking, 1987; Hawking, 1993], and the Gott causal violation proposal [Gott and Li, 1998] where the universe ‘creates itself’ and starts normal expansion in the domain without closed timelike lines.

Any of these may be combined with one or other proposals for

- (c) an effective ensemble of universes [Tegmark, 2003], realized either
 - (c1) in space-time regions that are part of either a larger entangled quantum entity, or are part of a single classical space-time, but are effectively disconnected from each other, or
 - (c2) in truly disconnected form.

All of these proposals however are strongly speculative, none being based solidly in well-founded and tested physics, and none being in any serious sense supported by observational evidence. They are all vast extrapolations from the known to the unknown. They may or may not be true. One thing is certain: they can’t all be true!

2.8 The concordance model

Observational support for the idea of expansion from a Hot Big Bang epoch is very strong, the linear magnitude-redshift relation for galaxies demonstrating the expansion,¹⁹ with source number counts and the existence of the blackbody CBR being strong evidence that there was indeed evolution from a hot early stage. Agreement between measured light element abundances and the theory of nucleosynthesis in the early universe confirms this interpretation. This basic theory is robust to critical probing. Much present activity attempts to link particle physics interactions during very early stages of the expansion of the universe to the creation of structures by gravitational instability much later, traces of the early seed fluctuations being accessible to us through present day CBR anisotropy patterns. Thus the present dominant cosmological paradigm is *a quantum gravity era of some kind followed by inflation; a hot big bang epoch; decoupling of matter and radiation; and then gravitational instability leading to formation of clusters of galaxies from the seed density perturbations that occur on the LSS.*

Together with supernova data, analysis of the CBR angular anisotropies and in particular their peaks gives a *concordance model* of this kind [Bennet et al., 2003; Tegmark, 2002; Tegmark *et al.*, 2004; Dodelson, 2003; Scott, 2005] that is then confirmed by the statistics of matter clustering [Eisenstein *et al.*, 2005a] together with observations of gravitational lensing and large-scale motions of matter [Silk, 2005]. This model is characterized by specific values for a set of cosmological parameters [Liddle, 2004], in particular

$$(37) \quad \Omega_{cdm0} \simeq 0.3, \Omega_{\Lambda 0} \simeq 0.7, T_{cbr0} = 2.75K, H_0 \simeq 65\text{km/sec/mpc}, t_0 \simeq 1.4 \times 10^{10}\text{years}.$$

Also $\Omega_{bar0} \simeq 0.044$ is the density of baryons, $\Omega_{vis0} \simeq 0.015$ that of luminous matter, and $\Omega_{\nu 0} \simeq 10^{-5}$ that of massless neutrinos, implying $\Omega_0 \simeq 0.3 + 0.7 \simeq 1$ in agreement with the inflationary prediction (34). The sign of k is uncertain, but if the combined evidence of all current observations is taken at face value it is positive, with $\Omega_0 = 1.02 \pm 0.02$ [Spergel *et al.*, 2003]. As noted above, there are some concerns firstly over age issues (see Sec. 2.3.2); secondly concerning the large angle CBR anisotropies (see Sec. 2.3.7); and thirdly regarding details of CDM structure formation at small scales (see Sec. 2.3.6); but none of these issues seems to be crucial at present.

2.8.1 Some misunderstandings

Despite its simplicity, there are some common misconceptions about the standard universe models (cf. [Lineweaver and Davis, 2005]) that can lead to philosophical misunderstandings.

Misconception 1: *The universe is expanding into something.* It is not, as it is all there is. It is just getting bigger, while always remaining all that is. One

¹⁹The alternative interpretation as gravitational redshifts in a static universe does not work because of the linearity of the observed redshift-distance relation [Ellis *et al.*, 1978].

should note here that a RW universe can be represented as a 4-dimensional curved spacetime expanding in a 5-dimensional flat embedding space time [Robertson, 1933]; however there is no necessity to view the 5-dimensional spacetime in this representation as physically real. Furthermore this embedding is no longer possible when we take perturbations into account; a 10 dimensional flat spacetime is needed for locally embedding a realistic (perturbed) universe model (and to do so globally requires many more dimensions, in general).

Misconception 2: *The universe expands from a specific point, which is the centre of the expansion.* All spatial points are equivalent in these universes, and the universe expands equally about all of them. Every fundamental observer sees exactly the same thing in an exact RW geometry. There is no centre to a FL universe.

Misconception 3: *Matter cannot recede from us faster than light.* It can, at an instant; two distantly separated fundamental observers in a surface $\{t = \text{const}\}$ can have a relative velocity greater than c if their spatial separation is large enough [Rothman and Ellis, 1993; Davis and Lineweaver, 2004]. No violation of special relativity is implied, as this is not a local velocity difference, and no information is transferred between distant galaxies moving apart at these speeds. For example, there is presently a sphere around us of matter receding from us at the speed of light;²⁰ matter beyond this sphere is moving away from us at a speed greater than the speed of light. The matter that emitted the CBR was moving away from us at a speed of about $61c$ when it did so [Rothman and Ellis, 1993].

Misconception 4: *The existence of a preferred RW frame (that in which the universe appears isotropic) contradicts relativity theory, which says all reference frames are equally good.* But this equivalence of frames is true for the equations rather than their solutions. Almost all particular solutions will have preferred world lines and surfaces; this is just a particular example of a *broken symmetry* — the occurrence of solutions of equations with less symmetries than the equations display. This feature is a key theme in modern physics [Brading and Castellani, 2006; Harvey, 2006].

Misconception 5: *The space sections are necessarily infinite if $k = 0$ or -1 .* This is only true if they have their ‘natural’ simply connected topology. If their topology is more complex (e.g. a 3-torus) they can be spatially finite [Ellis, 1971a; Lachièze *et al.*, 1995]. There are many ways this can happen; indeed if $k = -1$ there is an infinite number of possibilities.

Misconception 6: *Inflation implies spatial flatness ($k = 0 \Leftrightarrow \Omega_k = 1$) exactly.* There is nothing in inflationary theory which determines the sign of the spatial curvature. Inflationary universes are very nearly flat at late times; this is very different from being exactly flat (a condition which requires *infinite* fine tuning of initial conditions; if say the two millionth digit in the value of Ω_k is non-zero at

²⁰This sphere is not the same as the particle horizon, as is sometimes claimed (see [Rothman and Ellis, 1993]).

any time, then the universe is not spatially flat). Inflationary theory does not have the theoretical teeth required to imply that the universe has exactly flat spatial sections; hence a key issue for cosmology is observationally determining the sign of the spatial curvature, which is potentially dynamically important in both the very early universe [Ellis *et al.*, 2002b; Ellis *et al.*, 2002a] and the late universe (it determines if recollapse is possible, should the dark energy decay away).

2.8.2 Overall

Cosmology has changed from a speculative enterprise into a data-driven science that is part of standard physical theory [Barnett *et al.*, 1996]; a wealth of observations supports this dominant theory [Peebles *et al.*, 1991; Silk, 1997; Perkins, 2005]. Nevertheless some theoretical proposals are being made for the very early stages that have no observational support; and sometimes it may be impossible to ever obtain such support, both as regards the proposed physics and the geometry. Thus in some respects it remains a principle driven enterprise, with observation subordinate to theory.

We now explore the relation between cosmology and philosophy in terms of a series of *Theses* clustered around a set of major *Issues*. One can obtain a synoptic overview of the overall argument by simply considering the full set of *Issues* and *Theses*. They are summarized in the Table at the end.

3 ISSUE A: THE UNIQUENESS OF THE UNIVERSE.

The first and most fundamental issue is that there is only one Universe [Munitz, 1962; McCrea, 1960; Ellis, 1991]. This essential uniqueness of its object of study sets cosmology apart from all other sciences. In particular, the unique initial conditions that led to the particular state of the universe we see were somehow “set” by the time that physical laws as we know them started governing the evolution of both the universe and its contents, whenever that time may be. We cannot alter these unique initial conditions in any way — they are given to us as absolute and unchangeable, even though they are understood as contingent rather than necessary; that is, they could have been different while still being consistent with all known physical laws. The implications are that

Thesis A1: The universe itself cannot be subjected to physical experimentation. *We cannot re-run the universe with the same or altered conditions to see what would happen if they were different, so we cannot carry out scientific experiments on the universe itself.* Furthermore,

Thesis A2: The universe cannot be observationally compared with other universes. *We cannot compare the universe with any similar object, nor can we test our hypotheses about it by observations determining statistical properties of a known class of physically existing universes.*

Where this all becomes of observational relevance is in the idea of *cosmic variance* [Dodelson, 2003, pp. 241, 343]. The theory of structure formation in the early universe makes statistical predictions only (it cannot attempt to predict the specific structures that will actually be formed). Testing the theory compares our universe to a theoretical ensemble of universes, and declares a variance between what is measured in the actual universe and the expected properties based on the ensemble of models. If this variance is small enough, a deviation from expected values is pronounced as a statistical deviation, i.e. of no physical significance — we do not need to explain it any further; if it is large, it needs explanation. This is a key issue for example in the analysis of the CBR anisotropy observations [White *et al.*, 1993; Kamionkowski and Loeb, 1997]. The power spectrum of the CBR as measured by WMAP is less than expected at large angular scales (Sec. 2.3.7). One school of thought claims this is just a statistical fluctuation; another that it needs explanation, and might for example be evidence of a small universe [Luminet *et al.*, 2003; Luminet, 2005]. This debate arises because there is just one universe, and on large angular scales there are just a few measurements that can possibly be made (on small angular scales we can make many measurements and so this uncertainty becomes very small).

Consequent on **A1** and **A2**,

Thesis A3: The concept of ‘Laws of Physics’ that apply to only one object is questionable. *We cannot scientifically establish ‘laws of the universe’ that might apply to the class of all such objects, for we cannot test any such proposed law except in terms of being consistent with one object (the observed universe).*

This is insufficient: one observational point cannot establish the nature of a causal relation. Indeed the concept of a ‘law’ becomes doubtful when there is only one given object to which it applies [Munitz, 1962]. The basic idea of a physical law is that it applies to a set of objects all of which have the same invariant underlying behaviour (as defined by that law), despite the apparent variation in properties in specific instances, this variation resulting from varying initial conditions for the systems on which the law acts. This understanding is tested by physical experiments in which initial conditions for evolution of a set of similar systems are varied, and observations by which the statistical nature of a set of objects of the same broad kind is investigated. Neither is possible in the case of cosmology.

The laws of physics apply locally to the objects in the cosmos, and determine the evolution of the cosmos as a whole when locally applied everywhere with suitable initial/boundary conditions imposed (in the case of the RW models, via the Friedmann equation for example). Apart from this, we cannot establish higher-level effective laws that apply to all universes and determine their structure, as we can at all other levels of the hierarchy of complexity. All that we can do at this level of structure is observe and analyze the one unique object that exists. This is expressed by McCreia as follows: “When we speak of the other solutions of the equations of stellar structure, besides the one we are interested in at the moment, as representing systems that could exist, we mean that they could exist in the

universe as we know it. Clearly no such attitude is possible towards the universe itself" [McCrea, 1953].

Since the restriction of a global solution to a local neighborhood is also a solution, we have zillions of "mini-universe" on which to test the laws that control the local nature of the universe. But a mini-universe is not the universe itself; it is a small part of the whole. By examining these "mini-universes" and seeing if they are essentially the same everywhere, we can to some degree check firstly *that the laws of physics are the same everywhere in the universe* (a key feature of all cosmological analysis, cf. Sec. 7.1), and secondly *that the universe is spatially homogeneous* (this is discussed in depth below, see Sec. 4.2.2). But the latter feature is what has to be *explained* by a 'law of the universe'; verifying homogeneity does not explain why it is the case; this comes about because of specific initial conditions, which some suggest are due to hypothesized 'laws of the universe', applicable to the whole rather than to its parts. Finally,

Thesis A4: The concept of probability is problematic in the context of existence of only one object. *Problems arise in applying the idea of probability to cosmology as a whole — it is not clear that this makes much sense in the context of the existence of a single object which cannot be compared with any other existing object.*

But a concept of probability underlies much of modern argumentation in cosmology. Talk of 'fine tuning' for example is based on the use of probability (it is a way of saying something is improbable). This assumes both that things could have been different, and that we can assign probabilities to the set of unrealized possibilities in an invariant way. The issue here is to explain in what sense they could have been different with well-defined probabilities assigned to the different theoretical possibilities, if there is indeed only one universe with one set of initial conditions fixed somehow before physics came into being, or perhaps as physics came into being. We cannot scientifically establish laws of creation of the universe that might determine such initial conditions and resulting probabilities. If we use a Bayesian interpretation, which some suggest can be meaningfully applied to only one object [Garrett and Coles, 1993], the results depend on our 'prior knowledge', which in this case can be varied by changing our initial pre-physics assumptions. Related issues arise concerning the meaning of 'the wave function of the universe', at the heart of quantum cosmology. This wave function gives no unique prediction for any specific single universe.

Two comments on the above. First, *it is useful to distinguish between the experimental sciences — physics, chemistry, microbiology for example — on the one hand, and the historical and geographical sciences — astronomy, geology, evolutionary theory for example, on the other.* It is the former that are usually in mind in discussions of the scientific method. The understanding in these cases is that we observe and experiment on a class of identical or almost identical objects and establish their common behaviour. The problem then resides in just how identical those objects are. Quarks, protons, electrons, are all exactly identical to

each other, and so have exactly the same behaviour (indeed this feature underlies well-tested quantum statistics). All DNA molecules, frogs, human beings, and ecosystems are somewhat different from each other, but are similar enough nevertheless that the same broad descriptions and laws apply to them; if this were not so, then we would be wrong in claiming they belonged to the same class of objects in the first place. Water molecules, gases, solids, liquids are in an intermediate category — almost identical, certainly describable reliably by specific physical and chemical laws.

As regards the geographical and historical sciences, here one explicitly studies objects that are unique (the Rio Grande, the continent of Antarctica, the Solar System, the Andromeda galaxy, etc.) or events that have occurred only once (the origin of the Solar System, the evolution of life on Earth, the explosion of SN1987a, etc.). Because of this uniqueness, comment **A1** above applies in these cases also: we can only observe rather than experiment; the initial conditions that led to these unique objects or events cannot be altered or experimented with. However comment **A2** does not apply: at least in principle, there is a class of similar objects out there (other rivers, continents, planetary systems, galaxies, etc.) or similar events (the origin of other galaxies, the evolution of other planetary systems, the explosion of other supernovae, etc.) which we can observe and compare with our specific exemplar, also carrying out statistical analyses on many such cases to determine underlying patterns of regularity; and in this respect these topics differ from cosmology.

If we truly cannot carry out such analyses — that is, if **A2** applies as well in some particular case — then that subject partakes in this respect of the nature of cosmology. One may claim that *the dividing line here is that if we convince ourselves that some large-scale physical phenomenon essentially occurs only once in the entire universe, then it should be regarded as part of cosmology proper*; whereas if we are convinced it occurs in many places or times, even if we cannot observationally access them (e.g. we believe that planets evolved around many stars in other galaxies), then study of that class of objects or events can be distinguished from cosmology proper precisely because there is a class of them to study. The second comment is that some workers have tried to get around this set of problems by essentially *denying the uniqueness of the universe*. This is done by proposing the physical existence of ‘many universes’ to which concepts of probability can be properly applied (cf. Sec. 2.7.2), envisaged either as widely separated regions of a larger universe with very different properties in each region (as in chaotic inflation for example), as multiple realizations of quantum outcomes, or as an ensemble of completely disconnected universes — there is no physical connection whatever between them — in which all possibilities are realized. We return to this in Sec. 9.2.

4 ISSUE B: THE LARGE SCALE OF THE UNIVERSE IN SPACE AND TIME.

The problems arising from the uniqueness of the universe are compounded by its vast scale in both space and time, which poses major problems for observational cosmology. We therefore need to adduce various Principles in addition to the observations, in order to attain unique models: theory comes in as basis for interpreting observations.

4.1 Observations in a large scale universe

The distance to the nearest galaxy is about 10^6 light years, that is about 10^{24} cm., while the size of the earth is about 10^9 cm. The present size of the visible universe is about 10^{10} light years, that is about 10^{28} cm. This huge size relative to our own physical scale (about 10^2 cm) places major constraints on our ability to observe distant regions (and certainly prevents us experimenting with them). The uniqueness of cosmology in this respect is that it deals with this scale: the largest with which we can have causal or observational contact.

Thesis B1: Astronomical observations are confined to the past null cone, fading with distance. *We can effectively only observe the universe, considered on a cosmological scale, from one space-time event. Visual observations are possible only on our past light cone, so we are inevitably looking back into the past as we observe to greater distances. Uncertainty grows with distance and time.*

The vast scale of the universe implies we can effectively only view it from one spacetime event ('here and now') [Ellis, 1971a; Ellis, 1975]. If we were to move away from this spatial position at almost the speed of light for say 10,000 years, we would not succeed in leaving our own galaxy, much less in reaching another one; and if we were to start a long-term astronomical experiment that would store data for say 20,000 years and then analyze it, the time at which we observe the universe would be essentially unchanged (because its age is of the order of 10^{10} years: the extra time would make a negligible difference). This is quite unlike other geographic sciences: we can travel everywhere on earth and see what is there. The situation would be quite different if the universe were much smaller. Given its actual scale, such that we are now seeing galaxies whose present distance from us is about 10^9 light years, the effect is as if we were only able to observe the earth from the top of one mountain, and had to deduce its nature from those observations alone [Ellis, 1975].

Because we can only observe by means of particles — photons, massless neutrinos, gravitons — travelling to us at the speed of light, astronomical observations of distant sources and background radiation by telescopes operating at all wavelengths (optical, infrared, ultraviolet, radio, X-ray) are constrained to rays lying in our past light cone. These allow detailed observations (including visual pictures, spectral information, and polarization measurements) of matter as it intersects our past light cone. In observing distant regions, we can also aspire to use neutrino

and gravitational wave telescopes, and perhaps cosmic rays, also representing information coming to us at the speed of light or less. However all our detailed data about distant regions is gathered along our past light cone.

As a consequence, three interrelated problems occur in interpreting the astronomical observations. The first is that (because we can only view the universe from one point) *we only obtain a 2-dimensional projection on the sky of the 3-dimensional distribution of matter in the universe*. To reconstruct the real distribution, we need reliable distance measurements to the objects we see. However because of variation in the properties of sources, most are not reliable standard candles or standard size objects to use in calibrating distances, and in these cases we have to study statistical properties of classes of sources to estimate distances.

Second, *we necessarily see distant galaxies and other objects at earlier times in their history* (where their world lines intersect this past light cone).²¹ Thus cosmology is both a geographic and a historical science combined into one: we see distant sources at an earlier epoch, when their properties may have been different. As we are looking back in the past, source evolution must be taken into account; their properties at the time they emitted the light may be quite different from their properties now. We can only determine the distances of objects if we understand this evolution; but in practice it is one of the unknowns we have to try to determine (cf. Sec. 4.2.3).

Third, distant sources appear very small and very faint, both because of their physical distance, and because their light is highly redshifted (due to the expansion of the universe). Simply detecting them, let alone determining their characteristics, becomes rapidly more difficult with distance. Furthermore absorption by intervening matter can interfere with light from distant objects. The further back we look, the worse these problems become; thus our reliable knowledge of the universe decreases rapidly with distance [Ellis, 1975].

The situation is however improved by the availability of geological-type data [Hoyle, 1960]; that is, the present-day status of rocks, planets, star clusters, galaxies, and so on, which contains much information on the past history of the matter comprising those objects. Thus we can obtain detailed information on conditions near our past world-line in spacetime [Ellis, 1971a; Ellis, 1975] at very early times if we can interpret this data reliably, for example by relating theories of structure formation to statistical studies of source properties.

Thesis B2: ‘Geological’ type observations can probe the distant past of our past world line. *Physical and astrophysical observations tell us about conditions near matter world-lines in the far distant past. They can be used also to investigate the far distant past of more distant objects.*

This involves us in physical cosmology: namely the study of the evolution of structures in the universe, tested by comparison with astronomical observation. Particularly useful are measurements of the abundances of elements which resulted

²¹For example we see the Andromeda galaxy as it was two million years ago, long before humans existed on Earth [Silk, 2005].

from nucleosynthesis in the Hot Big Bang, giving us data about conditions long before decoupling (Sec. 2.2.2). If we can obtain adequate quality data of this kind for objects at high redshifts, we can use this to probe conditions very early on in their histories at some distance from our past worldline. Encouraging in this regard is the possibility of determination of element abundances at high redshift [Dodelson, 2003, pp. 11-12]; [Pettini, 1999]).

4.2 *Determining Spacetime Geometry: Observational Limits.*

The unique core business of observational cosmology is determining the large-scale geometry of everything there is, or at least of everything we can observe.

4.2.1 *Direct determination versus theory based approaches*

One can go about this in a direct manner: trying to determine the geometry of the universe directly from observations (assuming one has some understanding of the sources observed). The way this can be done (curiously known as the ‘inverse approach’) has been fully characterized [Kristian and Sachs, 1966; Ellis *et al.*, 1985]; indeed there is an interesting result here, namely

Observational Cosmology Theorem: *The data in principle available on our past null cone from astronomical observations is just necessary and sufficient to determine the space-time geometry on that null cone [Ellis *et al.*, 1985]. From this data one can in principle determine the space time in the past of the null cone and, if a no-interference conditions is assumed, to its future.*

However this is difficult to carry out both because of the problem of estimating distances for all observed sources, requiring a knowledge of the nature of the sources (Sec. 4.2.3),²² and because of the serious difficulty in obtaining some of the needed data (which include apparent distortions of all distant objects, and the transverse velocities of all observed matter). The further we observe down the past light cone, the larger the uncertainty becomes. This direct observational approach, where no prior model is assumed for the space-time geometry, has been pursued to some degree (and in essence underlies for example the observational studies that discovered large-scale structure such as the great walls and voids). Nevertheless it is not widely adopted as an overall approach to cosmology, both because of these observational difficulties, but also because it has little explanatory value; it just tells us what the geometry and matter distribution is, but not why it is of that nature.

The usual option in cosmology proper is rather to use a theory-based approach: we *a priori* assume a model based on a space-time geometry with high symmetry (usually a FL model, see Sec. 2.1), and then determine its essential free parameters from comparison of theoretical relations with astronomical observations (Sec. 2.3.3). Detailed observations of the matter distribution and large-scale velocities

²²The link between observations and models always requires some theory, and is never direct.

as well as CBR anisotropies then help us determine deviations from the exact model, both statistically (an astrophysical description [Dodelson, 2003]) and in detail (an astronomical description [Ellis and Stoeger, 1987]).

4.2.2 Indirect determination: justifying a Friedmann-Lemaître geometry

The standard models of cosmology are the Friedmann-Lemaître (FL) family of universe models that are exactly spatially homogeneous and isotropic everywhere (Sec. 2.1). They are easy to understand, and have tremendous explanatory power; furthermore their major physical predictions (the existence of blackbody CBR and specific light element production in the early universe) seem confirmed. The issue is, to what degree does observational data uniquely indicate these universe models for the expanding universe geometry? Here one is assuming a large enough averaging scale for spatial homogeneity to be valid; this scale should be explicitly indicated [Ellis, 1984] (it is about 100 Mpc at present [Dodelson, 2003]).²³ These are the background models for cosmology; perturbed FL models then characterize the nature of deviations from the exact FL geometry that are expected on smaller scales (Sec. 2.5.2).

The key feature here is the observed *isotropy about our location* (Sec. 2.3.1). Considered on a large enough angular scale, astronomical observations are very nearly isotropic about us, both as regards source observations and background radiation; indeed the latter is spectacularly isotropic, better than one part in 10^4 after a dipole anisotropy, understood as resulting from our motion relative to the rest frame of the universe, has been removed [Partridge, 1995]. Because this applies to all observations (in particular, there are not major observed matter concentrations in some other universe region), this establishes that in the observable region of the universe, to high accuracy *both the space-time structure and the matter distribution are isotropic about us*. We can easily construct exact spherically symmetric universe models [Bondi, 1947; Ellis and van Elst, 1999a], as indicated by these observations. In general they will be spatially inhomogeneous, with our Galaxy located at or near the centre; this is currently a philosophically unpopular proposal, but is certainly possible. The question is whether we can give convincing observational evidence for spatial homogeneity in addition to the spherical symmetry. Various arguments are used for this purpose.

- (a) *The cosmological principle* [Bondi, 1960; Weinberg, 1972]: Just assume spatial homogeneity because it is the simplest case and you don't need anything more complex on the basis of current data. We simply adopt a philosophical principle as the basis of argument. This is essentially an *a priori prescription for initial conditions for the universe* (a universe that initially has a RW geometry will have that geometry at later times, because symmetries of

²³There exist *hierarchical models* where neither the fluid approximation nor homogeneity is ever attained at any scale because of their fractal nature [de Vaucouleurs, 1970]. The regularity of the observed galactic motions, as evidenced by the (m, z) relations, speaks against these models, as do large-scale observations of the matter distribution [Peebles, 1993a].

the initial data are preserved by the Einstein equations [Hawking and Ellis, 1973]); but it is not usually expressed that way.

- (b) *FL observational relations*: If we could show that the source observational relations had the unique FL form (26, 28) as a function of distance, this would establish spatial homogeneity in addition to the isotropy, and hence a RW geometry [Ellis *et al.*, 1985]. This is essentially what is done for example in using number counts towards establishing spatial homogeneity [Hubble, 1936]. However because of Thesis **B1** above, the observational problems mentioned earlier — specifically, unknown source evolution — prevent us from carrying this through: we cannot measure distances reliably enough. Astrophysical cosmology could resolve this in principle, but is unable to do so in practice. Indeed the actual situation is the inverse: *taking radio-source number-count data at its face value, without allowing for source evolution, contradicts a RW geometry.*

In the face of this, the usual procedure is to assume that spatial homogeneity is known in some other way, and deduce the source evolution required to make the observations compatible with this geometric assumption [Ellis, 1975]. It is always possible to find a source evolution that will achieve this [Mustapha *et al.*, 1998]. Thus attempts to observationally prove spatial homogeneity this way fail; indeed an alternative interpretation would be that this data is evidence of spatial inhomogeneity, i.e. that we live in a spherically symmetric inhomogeneous universe where we are situated somewhere near the centre, with the cosmological redshift being partly gravitational, cf. [Ellis *et al.*, 1978] (and conceivably with a contribution to the CBR dipole from this inhomogeneity if we are a bit off-centre). Similarly the supernova data usually understood as implying the existence of a cosmological constant (Sec. 2.3.5) could also be interpreted in this way as evidence of inhomogeneity, without the need for ‘dark energy’. Most people regard such proposals as very unappealing — but that does not prove they are incorrect.

- (c) *Physical arguments*: One can claim that physical processes such as inflation (Sec. 2.6) make the existence of almost-RW regions highly likely, indeed much more probable than spherically symmetric inhomogeneous regions. This is a viable argument, but we must be clear what is happening here — we are replacing an observational test by a theoretical argument based on a physical process that may or may not have happened (for there is no definitive observational proof that inflation indeed took place). It is strongly bolstered because predictions for the detailed pattern of CBR anisotropy on small scales [Hu and Sugiyama, 1995b], based on the inflationary universe theory, have been confirmed [Perkins, 2005]; but that argument will only become rigorous if it is shown that spherically symmetric inhomogeneous models (with or without inflation) cannot produce similar patterns of anisotropy. But they probably can, because the acoustic oscillations that lead to the characteristic predicted anisotropy patterns in fact take place after inflation,

and can equally happen if suitable initial conditions occur without a previous inflationary phase.

What about alternative observational routes? Another proposal is,

- (d) *Uniform thermal histories*: the idea is to use the uniformity in the nature of the objects we see in the sky (we see the same types of galaxy at large distances, for example) to deduce they must have all undergone essentially the same thermal history, and then to prove from this homogeneity of thermal histories that the universe must be spatially homogeneous. For example, observations showing that element abundances at high redshift in many directions are the same as locally, are very useful in constraining inhomogeneity by showing that conditions in the very early universe at the time of nucleosynthesis must have been the same at distant locations in these directions [82]. However turning this idea into a proper test of homogeneity has not succeeded so far: indeed it is not clear if this can be done, because some (rather special) counter-examples to this conjecture have been found [Bonnor and Ellis, 1986]. Nevertheless the approach could be used to give evidence against spatial homogeneity: for example, if element abundances were measured to be different at high redshifts in any direction [Pettini, 1999; Sigurdson and Furlanetto, 2005], or if ages of distant objects were incompatible with local age estimates [Jain and Dev, 2005].

Finally the argument for spatial homogeneity that is most generally accepted:

- (e) *Isotropy everywhere*: If all observers see an isotropic universe, then spatial homogeneity follows [Walker, 1944; Ehlers, 1993; Ellis, 1971a]; indeed homogeneity follows if only three spatially separated observers see isotropy. Now we cannot observe the universe from any other point, so we cannot observationally establish that far distant observers see an isotropic universe. Hence the standard argument is to assume a *Copernican Principle*: that we are not privileged observers. This is plausible in that all observable regions of the universe look alike: we see no major changes in conditions anywhere we look. Combined with the isotropy we see about ourselves, this implies that *all observers see an isotropic universe*, and this establishes a RW geometry [Walker, 1944; Ellis, 1971a; Hawking and Ellis, 1973]. This result holds if we assume isotropy of *all* observations; a powerful enhancement was proved by Ehlers, Geren, and Sachs [Ehlers *et al.*, 1968; Hawking and Ellis, 1973], who showed that if one simply assumes isotropy of freely-propagating radiation about each observer in an expanding universe domain,²⁴ the result follows from the Einstein and Liouville equations; that is,

²⁴This result does not hold in a static universe, for then the radiation temperature depends only on the potential difference between the emitter and observer, hence the radiation is isotropic everywhere even if the universe inhomogeneous, cf. [Ellis *et al.*, 1978].

EGS Theorem: *Exact isotropy of the CBR for every geodesically moving fundamental observer at each point in an expanding universe domain U implies an exact RW geometry in U .*

Thus we may establish spatial homogeneity by assuming a weak Copernican principle: we are not in a privileged position where the CBR just happens to be highly isotropic by chance; hence all comoving observers may be assumed to measure highly isotropic CBR, and the result follows. This is currently the most persuasive observationally-based argument we have for spatial homogeneity.

A problem is that it is an exact result, assuming exact isotropy of the CBR; is the result stable? Indeed it is: *almost-isotropy of freely-propagating CBR for an expanding family of geodesically-moving fundamental observers everywhere in some region proves the universe geometry is almost-RW in that region* [Stoeger *et al.*, 1995]. Thus the result applies to the real universe — provided we make the Copernican assumption that all other observers, like us, see almost isotropic CBR. And that is the best we can do at present. Weak tests of the isotropy of the CBR at other spacetime points come from the Sunyaev-Zel'dovich effect [Goodman, 1995] and from CBR polarization measurements [Kamionkowski and Loeb, 1997], giving broad support to this line of argument but not enough to give good limits on spatial inhomogeneity.

The observational situation is clear:

Thesis B3: Establishing a Robertson-Walker geometry for the universe relies on plausible philosophical assumptions. *The deduction of spatial homogeneity follows not directly from astronomical data, but because we add to the observations a philosophical principle that is plausible but untestable.*

The purpose of the above analysis is not to seriously support the view that the universe is spherically symmetric and inhomogeneous, as is allowed by the observations, but rather to show clearly the nature of the best observationally-based argument by which we can (quite reasonably) justify the assumption of spatial homogeneity.

Accepting this argument, the further question is, *in which spacetime regions does it establish a RW-like geometry?* The CBR we detect probes the state of the universe from the time of decoupling of matter and radiation (at a redshift of about 1100) to the present day, within the visual horizon. The argument from CBR isotropy can legitimately be applied for that epoch. However, it does not necessarily imply isotropy of the universe at much earlier or much later times, because there are spatially homogeneous anisotropic perturbation modes that are unstable in both directions of time; and they will occur in a generic situation. Indeed, if one examines the Bianchi (spatially homogeneous but anisotropic) universes, using the powerful tools of dynamical systems theory, one can show that *intermediate isotropisation* can occur [Wainwright and Ellis, 1996; Wainwright *et al.*, 1998]: *despite being highly anisotropic at very early and very late times, such models can mimic a RW geometry arbitrarily closely for an arbitrarily long time*, and hence can reproduce within the errors any set of FL-like

observations. We can obtain strong limits on the present-day strengths of these anisotropic modes from CBR anisotropy measurements and from data on element abundances, the latter being a powerful probe because (being of the ‘geological’ kind) they can test conditions at the time of element formation, long before decoupling. But however low these observational limits, anisotropic modes can dominate at even earlier times as well as at late times (long after the present). If inflation took place, this conclusion is reinforced: it washes out any information about very early universe anisotropies and inhomogeneities in a very efficient way.

As well as this time limitation on when we can regard homogeneity as established, there are major spatial limitations. The above argument does not apply far outside the visual horizon, for we have no reason to believe the CBR is highly isotropic there. Indeed if chaotic inflation is correct, conditions there are not the same.

4.2.3 Determining the RW parameters

Given that a RW geometry is a good description of the observable universe on a large scale, the further issue is what are the best-fit parameters that characterize it, selecting the specific universe we observe from the family of all FL models (Sec. 2.1). Important observational issues are:

- Determining the Hubble parameter H_0 , which sets the overall scale of the observed universe region.
- Determining the trio of the density parameter Ω_0 , deceleration parameter q_0 , and cosmological constant Λ (or equivalently the density parameter Ω_Λ), which are the major defining characteristics of a specific FL model. The CBR data, supernova observations, deep number counts, source covariance functions, velocity measurements, and gravitational lensing observations can determine these quantities.
- Determining the sign of the curvature k , showing whether the universe has closed spatial sections and also whether it is possible for it to recollapse in the future or not. Analyses of the observations should always attempt to determine this sign, and not assume that $k = 0$ (as is often done) [Wright, 2006].
- Various parameters are used to characterize the nature of dark matter (Sec. 2.3.6) and dark energy (Sec. 2.3.5). As their dynamics is unknown, these too have to be determined observationally.

We only obtain good estimates of these quantities by the observational relationships characterized above (Sec. 2.3.3) using statistical analysis of the classes of objects we observe. Problems arise because of our lack of adequate theories of their historical development.

Thesis B4: Interpreting cosmological observations depends on astrophysical understanding. *Observational analysis depends on assessing a variety of auxiliary functions characterizing the sources observed and the observations made. These introduce further parameters that have to be observationally or theoretically determined, allowing considerable freedom in fitting specific models to the observations. Physical cosmology aims to characterize perturbed FL models (which account for structure formation) rather than just the background exactly smooth FL models; this introduces further parameters to be determined.*

It is useful here to distinguish between methods aimed at determining the properties of the background (zeroth order) FL model directly, and those aimed at determining properties of the perturbations of these models [Tegmark, 2002]. Methods for determining the parameters of the background model (Sec. 2.1) depend on assuming properties of the distance indicators used (galaxies, radio sources, etc.). They will have their own properties (brightness profiles, luminosities, physical sizes, spectra, etc.) and dynamical evolution; but these are often not well understood, and will have to be represented in a parametric way (e.g. by parameters describing luminosity evolution). In each case we end up assuming important aspects of the astrophysics and evolutionary histories of the objects observed, which are not part of the cosmological model proper. The statistical properties of the sources observed are also characterized by parametrized functions (e.g. the luminosity function characterizing the numbers of galaxies in each luminosity class) that have to be known in order to analyze the observations. This situation is an example of Lakatos' view of how scientific programmes work, with a belt of auxiliary hypotheses interposing between the core theoretical proposal and the data used to test it [Lakatos, 1980]. This makes the analysis rather model-dependent, where the models are only indirectly related to the background model — their explanation is the aim of astrophysics rather than cosmology. Thus if observational results disagree with a particular cosmological model, one can always claim it is the understanding of the auxiliary hypotheses that is at fault rather than the model being proposed [Lakatos, 1980].

By contrast, many of the methods of estimating Ω_0 (and to some degree Λ) depend on studying the growth and nature of inhomogeneities in the universe, that is they investigate perturbed FL models (Sec. 2.5.2), whose properties of course depend on the background model, but introduce a whole set of further functions and parameters describing the perturbations [Dodelson, 2003], for example the angular correlation function for matter (or its Fourier transform, the 2-dimensional power spectrum), the power spectrum of density fluctuations [Tegmark, 2002], red-shift space correlation functions [Peebles, 1993a; Eisenstein *et al.*, 2005a], and correlation function for velocities [Dodelson, 2003]. Associated parameters include a scalar *spectral index* (characterizing the spectrum of physical sizes of inhomogeneities), the *bias parameter* b (expressing how galaxy formation is biased towards density peaks in the inhomogeneities [Dodelson, 2003, p. 280]) and the *initial fluctuation magnitudes* Q (the seeds for structure formation). Determining these parameters is part of the task of cosmology proper: to fully characterize the

perturbed cosmological model, *we aim to determine both the background parameters and the quantities describing the perturbations*. Model selection then depends on the parameters used to describe them — what is assumed known, and what is to be determined [Liddle, 2004; Scott, 2005]. For example, standard inflationary theory predicts a scale-invariant spectrum of Gaussian perturbations; do we test that assumption, or take it for granted? This comes up in the issue of what ‘priors’ are assumed when conducting statistical tests.

4.2.4 Consistency tests

A key question for cosmology is *what kinds of observations provide critical tests of the standard FL models*. If there were no observations that could disprove them, the subject would be of questionable scientific status. An important such test is obtaining *estimates of the age of the universe* t_0 , which is dependent on H_0 , Ω_0 , and Λ , and comparing them with estimates of the ages of objects in the universe (determined on astrophysical grounds):

Thesis B5: A crucial observational test for cosmology is that the age of the universe must be greater than the ages of stars. *The tension between the age of the universe and ages of stars is one area where the standard models are vulnerable to being shown to be inconsistent, hence the vital need to establish reliable distance scales, basic to estimates of both H_0 and the ages of stars, and good limits on Λ . Other consistency tests help confirm the standard model and consolidate cosmology’s standing as an empirical science.*

At present this age issue is acceptable for local objects, because of a recent revision of our distance scale estimates [Harris *et al.*, 1998], assisted by data that Λ is positive [Perlmutter *et al.*, 1998]; but continued vigilance is needed on this front, particularly as there are indications of problems for high redshift objects [Jain and Dev, 2005]. If this ever became serious we might have to resort to spherically symmetric inhomogeneous models rather than spatially homogeneous models, with the ‘bang time’ (characterizing the start of the universe) dependent on distance from us [Mustapha *et al.*, 1998].

Note that this issue is crucially unlike the case of the large angle CBR anisotropies (Sec. 2.3.7): *the low CBR anisotropies at large angular scales can as a last resort be dismissed as a statistical fluke; the age issue cannot*. It is to do with the internal consistency of individual cosmological models, not with probabilities. Thus it is a plus for cosmology that the age issue exists. Other consistency tests include

- Showing that the CBR temperature T_{cbr} varies with redshift according to $T_{cbr} = 2.75(1+z)$ [Meyer, 1994];
- Confirming that helium abundances are consistent with a primordial value of 25% at large distances (high redshifts) in all directions [Dodelson, 2003, pp. 11-12]; also [Pettini, 1999; Sigurdson and Furlanetto, 2005]; and

- Checking that there is a 2% number count dipole parallel to the CBR dipole for all cosmological sources [Ellis and Baldwin, 1984].

4.3 *The hidden universe*

If we do not live in a small universe (Sec. 4.3.1), the further essential point is that the region of the universe we can observe is restricted, firstly because we cannot see to earlier times than the LSS (the universe was opaque before then (see Sec. 2.2)), and secondly because a finite time has elapsed since the universe became transparent to radiation, and light can only have travelled a finite distance in that time. As no signal can travel to us faster than light, we cannot receive any information from galaxies more distant than our visual horizon [Ellis and Stoeger, 1988]. The most distant matter we can observe is that which emitted the CBR (Sec. 2.4.2).

Thesis B6: Observational horizons limit our ability to observationally determine the very large scale geometry of the universe. *We can only see back to the time of decoupling of matter and radiation, and so have no direct information about earlier times; and unless we live in a ‘small universe’, most of the matter in the universe is hidden behind the visual horizon. Conjectures as to its geometry on larger scales cannot be observationally tested. The situation is completely different in the small universe case: then we can see everything there is in the universe, including our own galaxy at earlier times.*

The key point here is that unless we live in a small universe, *the universe itself is much bigger than the observable universe.* There are many galaxies — perhaps an infinite number — at a greater distance than the horizon, that we cannot observe by any electromagnetic radiation. Furthermore no causal influence can reach us from matter more distant than our particle horizon — the distance light can have travelled since the creation of the universe, so this is the furthest matter with which we can have had any causal connection [Rindler, 1956; Hawking and Ellis, 1973; Tipler *et al.*, 1980]. We can hope to obtain information on matter lying between the visual horizon and the particle horizon by neutrino or gravitational radiation observatories; but we can obtain no reliable information whatever about what lies beyond the particle horizon. We can in principle feel the gravitational effect of matter beyond the horizon because of the force it exerts (for example, matter beyond the horizon may influence velocities of matter within the horizon, even though we cannot see it). This is possible because of the constraint equations of general relativity theory, which are in effect instantaneous equations valid on spacelike surfaces.²⁵ However we cannot uniquely decode that signal to determine what matter distribution outside the horizon caused it: a particular velocity field might be caused by a relatively small mass near the horizon, or a much larger

²⁵They are valid at any late time in a solution of the EFE because they were valid initially — the initial data must satisfy constraint equations — and once they are satisfied, the constraints are preserved by the dynamic field equations.

mass much further away [Ellis and Sciama, 1972]. Claims about what conditions are like on very large scales — that is, much bigger than the Hubble scale — are unverifiable [Ellis, 1975], for we have no observational evidence as to what conditions are like far beyond the visual horizon. The situation is like that of an ant surveying the world from the top of a sand dune in the Sahara desert. Her world model will be a world composed only of sand dunes — despite the existence of cities, oceans, forests, tundra, mountains, and so on beyond her horizon.

It is commonly stated that if we live in a low-density universe and the cosmological constant vanishes, the universe has infinite spatial sections. However this deduction only applies if firstly the RW-like nature of the universe within the past light cone continues to be true indefinitely far outside it, and secondly the space sections have their ‘natural’ simply-connected topology — and there is no way we can obtain observational evidence that these conditions are both true. In contrast to this, in chaotic inflationary models (Sec. 2.6), it is a definite prediction that the universe will not be like a RW geometry on a very large scale — rather it will consist of many RW-like domains, each with different parameter values, separated from each other by highly inhomogeneous regions outside our visual horizon [Linde, 1990], the whole forming a fractal-like structure. This prediction is just as untestable as the previously prevalent assumption (based on a Cosmological Principle) that the universe is RW-like on such scales [Bondi, 1960; Weinberg, 1972]. Neither can be observationally confirmed or denied. The same issue arises in an even more extreme form in relation to the idea of a multiverse. We return to this below, see Sec. 9.2.

4.3.1 *Small universes*

There is one case where this kind of spatial observational limit does not obtain. This is when a *Small Universe* occurs, that is, a universe which closes up on itself spatially for topological reasons [Ellis, 1971b], and does so on such a small scale that we have seen right round the universe since the time of decoupling. Then we can see all the matter that exists, with multiple images of many objects occurring [Ellis and Schreiber, 1986]. This possibility is observationally testable by examining source statistics, by observation of low power in the large angle CBR anisotropies, and by detecting identical temperature variation on various circles in the CBR sky [Lachièze *et al.*, 1995]. There are weak hints in the observed CBR anisotropies (the lack of power on large angular scales) that this could actually be the case [Luminet *et al.*, 2003; Luminet, 2005], but this is not solidly confirmed. Checking if the universe is a small universe or not is an important task; the nature of our observational relationship to the universe is fundamentally different if it is true [Ellis and Schreiber, 1986].

4.4 *The observed universe*

The observable part of the universe (i.e. back to the visual horizon) is strictly limited, and we have already seen most of it. We can only observe distant objects

by electromagnetic radiation at all wavelengths, by neutrinos, and by gravitational waves. We already have very complete broad coverage of the entire sky by electromagnetic observations at all wavelengths right back to the surface of last scattering, which is the limit of what will ever be observable by electromagnetic radiation. Detailed observations (such as the Hubble Deep Field) are available for restricted domains in angle and depth. Detailed observations at suitable wavelengths are beginning to discern what lies behind the Milky Way, which tends to obscure a substantial fraction of the sky. It is unlikely there are many new astronomical phenomena undiscovered in this observable region, although it will be crucial determining more detailed features of the phenomena we have already discovered (e.g. the nature of dark matter and dark energy).

Thesis B7: We have made great progress towards observational completeness. *We have already seen most of the part of the universe that is observable by electromagnetic radiation. It is plausible that not many new astronomical phenomena remain to be discovered by us observationally; we will determine more details (so understanding more about what we have seen) and see more objects, but not discover many new kinds of things.*

Indeed Harwit [1984] has used the multiplicity of discovery of specific astronomical phenomena to estimate how many new essentially different such phenomena there are still waiting to be discovered.

Neutrinos and gravitational waves will in principle allow us to peer back to much earlier times (the time of neutrino decoupling and the quantum gravity era respectively), but are much harder to observe at all, let alone in useful directional detail. Nevertheless the latter has the potential to open up to us access to eras quite unobservable in any other way. Maybe they will give us unexpected information on processes in the very early universe which would count as new features of physical cosmology.

5 ISSUE C: THE UNBOUND ENERGIES IN THE EARLY UNIVERSE

The analogous problems for physical cosmology arise because energies occurring in the Hot Big Bang early universe phase (Sec. 2.2) are essentially unbounded, so the highest energies we can attain in particle accelerators cannot reach the levels relevant to very early times. The uniqueness of cosmology in this regard is that it is the only science contemplating spacetime regions that have experienced such high energies, and with which we are in intimate causal contact despite the huge timescales involved — indeed events at those early times determined much of what we see around us today. The nuclear reactions underlying nucleosynthesis are well understood, and their cross-sections reasonably well-known; the processes of baryosynthesis and quark-gluon recombination are reasonably understood and are on the border of being testable; but physical processes relevant at earlier times are inaccessible to testing by laboratory or accelerator-based experiment. The *Physics Horizon* by definition separates those aspects of physics we can hope to test by

high-energy experiments on Earth or in the Solar System, from those where it is reasonable to expect no such test will ever be possible:

Thesis C1: The Physics Horizon limits our knowledge of physics relevant to the very early universe. *We cannot experimentally test much of the physics that is important in the very early universe because we cannot attain the required energies in accelerators on Earth. We have to extrapolate from known physics to the unknown and then test the implications; to do this, we assume some specific features of known lower energy physics are the true key to how things are at higher energies. We cannot experimentally test if we have got it right.*

Note that this is independent of the issue of setting of initial conditions for the universe, considered below, see Sec. 6.2: the problem arises after the initial conditions have been set and the universe is running according to invariable physical laws. We cannot be confident of the validity of the physics we presuppose then. Rather than using known physics to predict the evolution of the universe, *we end up testing proposals for this physics by exploring their implications in the early universe*, which is the only ‘laboratory’ where we can test some of our ideas regarding fundamental physics at the highest energies [Yoshimura, 1988]; this is particularly true in the case of quantum gravity proposals. The problem is we cannot simultaneously do this and also carry out the aim of physical cosmology, namely predicting the evolution of the early universe from known physical theory.

Our understanding of physics at those times has of necessity to be based on extrapolation of known physics way beyond the circumstances in which it can be tested. The trick is to identify which features are the key to use in that extrapolation: for example, variational principles, broken symmetries and phase changes, duality invariance, entropy limits are candidates. If we confirm our guesses for the relevant physics by their satisfactory implications for the early universe, tested in some suitable way, then this is impressive progress; but if this is the *only* way we can test the proposed physics, the situation is problematic. If the hypothesis solves only the specific issues it was designed to solve in the early universe and nothing else, then in fact it has little explanatory power, rather it is just an alternative (perhaps theoretically preferable) description of the known situation. One obtains positive observational support for a particular proposal for the relevant physics only if it predicts multiple confirmed outcomes (rather than just one), for example predicting particles that are then confirmed to exist in a laboratory, so that a single hypothesis simultaneously solves several different observational issues. Some of the options may be preferred to others on various theoretical grounds; but one must distinguish this from their having observational support. They lack physical power if they have no other testable consequences. A particular example is the inflationary universe proposal (Sec. 2.6): the supposed inflaton field underlying an inflationary era of rapid expansion in the early universe [Guth, 1981; Gibbons *et al.*, 1983; Kolb and Turner, 1990; Guth, 1997] has not been identified, much less shown to exist by any laboratory experiment. Because this field ϕ is unknown, one can assign it an arbitrary potential $V(\phi)$, this arbitrariness reflecting

our inability to experimentally determine the relevant behaviour. It can be shown that virtually any desired scale evolution $S(t)$ of the universe can be attained by suitable choice of this potential [Ellis and Madsen, 1991]; and also almost any desired perturbation spectrum can be obtained by a (possibly different) suitable choice [Lidsey *et al.*, 1997]. Indeed in each case one can run the mathematics backwards to determine the required potential $V(\phi)$ from the desired outcome (Sec. 9.3.1 below). The mathematical existence of such a theoretical potential of the desired form for cosmological purposes does not by itself prove a particle or field exists with that effective potential.

Thesis C2: The unknown nature of the inflaton means that inflationary universe proposals are incomplete. *The promise of inflationary theory in terms of relating cosmology to particle physics has not been realized. This will only be the case when the nature of the inflaton has been pinned down to a specific field that experiment confirms or particle physics requires to exist.*

The very impressive achievement of inflation is that the predicted CBR anisotropy spectrum is verified and agrees with the matter power spectrum [Eisenstein *et al.*, 2005a]; but that prediction depends only on the physics from the era of tight coupling of matter and radiation to the present day, given a suitable initial fluctuation spectrum in the early universe, rather than on the specific hypothesis of an inflationary origin for that spectrum. The true clincher would be if properties of an inflationary field were predicted from the cosmology side and then confirmed in the laboratory; indeed that would count as one of the great feats of theoretical physics. This may not happen however because of the experimental problems focused on here, arising because we cannot reproduce on Earth all the conditions relevant to very early cosmology.

One key application where this issue becomes significant is in respect of the chaotic inflation theory (Sec. 2.6). As remarked above, see Sec. 4.3, its geometric predictions are observationally unverifiable. It would nevertheless be a good physical prediction if it was a more or less inevitable outcome of known and tested underlying physics. However this is not the case: the proposed underlying physics is not experimentally tested, indeed it is not even uniquely defined or associated with any specific known physical particle or field. The claim that it inevitably follows from string theory [Susskind, 2005] suffers from the problem that string theory is not a well-defined or tested part of physics.

6 ISSUE D: EXPLAINING THE UNIVERSE — THE QUESTION OF ORIGINS.

This is the unique core business of physical cosmology: explaining both why the universe has come into existence and evolved to the present very high-symmetry FL geometry on large scales, and how structures come into existence on smaller scales.

6.1 *Start to the universe*

Did a start to the universe happen? If so, what was its nature? This has been discussed above (Sec. 2.7.2), and the issue is unresolved. The major related question is whether the process of expansion only happens once in the life of the Universe, or occurs repeatedly. The first option is the standard model, where the entire evolution of the Universe is a once-off affair, with all the objects we see, and indeed the Universe itself, being transient objects that will burn out like dead fireworks after a firework display. In this case everything that ever happens occurs during one expansion phase of the Universe (possibly followed by one collapse phase, which could occur if $k = +1$ and the present ‘dark energy’ field dies away in the future). This evolution might have a singular start at a space-time singularity; a beginning where the nature of time changes character; a non-singular bounce from a single previous collapse phase; or a start from a non-singular static initial state [Mulryne *et al.*, 2005]. An alternative is that many such phases have occurred in the past, and many more will occur in the future; the Universe is a *Phoenix Universe* [Dicke and Peebles, 1979], new expansion phases repeatedly arising from the ashes of the old. While the idea of one or more bounces is an old one [Tolman, 1934], actual mechanisms that might allow this bounce behaviour have not yet been elucidated in a fully satisfactory way. A variant is the chaotic inflation idea of new expanding universe regions arising from vacuum fluctuations in old expanding regions, leading to a universe that has a fractal-like structure at the largest scales, with many expanding regions with different properties emerging out of each other in a universe that lasts forever (Sec. 2.6).

As discussed above, see Sec. 2.7.1, it is possible (if the universe has positive spatial curvature) that the quantum gravity domain can be avoided and there was no start to the universe; however this probably requires special initial conditions [Ellis and Maartens, 2004]. If a quantum gravity domain indeed occurred, we cannot come to a definite conclusion about whether there was a creation event or not because we do not know the nature of quantum gravity, nor how to reliably apply it in the cosmological context where the issue of initial conditions arises. Loop quantum gravity suggests the universe may be singularity-free [Bojowald, 2001], with bounces or a non-singular start, but that theory is unconfirmed. Tested physics cannot give a decisive answer; it is possible that *testable* physics also cannot do so.

Thesis D1: An initial singularity may or may not have occurred. *A start to the universe may have occurred a finite time ago, but a variety of alternatives are conceivable: eternal universes, or universes where time as we know it came into existence in one or another way. We do not know which actually happened, although quantum gravity ideas suggest a singularity might be avoided.*

This is a key issue in terms of the nature of the universe: a space-time singularity is a dramatic affair, where the universe (space, time, matter) has a beginning and all of physics breaks down and so the ability to understand what happens on a scientific basis comes to an end. However eternal existence is also problematic,

leading for instance to the idea of Poincaré’s eternal return: everything that ever happened will recur an infinite number of times in the future and has already occurred an infinite number of times in the past [Barrow and Tipler, 1984]. This is typical of the problems associated with the idea of infinity (discussed further below, see Sec. 9.3.2). *It is not clear in the end which is philosophically preferable: a singularity or eternal existence.* That decision will depend on what criteria of desirability one uses (such criteria are discussed below, see Sec. 8.1).

6.2 *The issue of initial conditions*

While occurrence of an initial singularity is striking in that it is a start to physics and spacetime as well as matter, whether it occurred or not is in a sense irrelevant to the key issue of what determined the nature of the universe:

Thesis D2: Testable physics cannot explain the initial state and hence specific nature of the universe. *A choice between different contingent possibilities has somehow occurred; the fundamental issue is what underlies this choice. Why does the universe have one specific form rather than another, when other forms consistent with physical laws seem perfectly possible? The reasons underlying the choice between different contingent possibilities for the universe (why one occurred rather than another) cannot be explored scientifically. It is an issue to be examined through philosophy or metaphysics.*

Even if a literal creation does not take place, as is the case in various of the present proposals, this does not resolve the underlying issue of what determined why the universe is the way it is, given that it could presumably have been otherwise. If the proposal is evolution from a previous eternal state — Minkowski space for example — then why did that come into existence, and why did the universe expand as a bubble from that vacuum start when it did, rather than at some previous time in the pre-existent eternity? Whenever it started, it could have started before! Some attempts involve avoiding a true beginning by going back to some form of eternal or cyclic initial state, for example Tolman’s series of expansion and collapse cycles [Tolman, 1934], proposals for creation of the universe as a bubble formed in a flat space-time [Tryon, 1973], Linde’s eternal chaotic inflation [Linde, 1990], Veneziano’s re-expansion from a previous collapse phase [Ghosh *et al.*, 1998], the ekpyrotic universe proposal [Khoury *et al.*, 2001], and theories involving foundational limits on information through a “holographic principle” [Susskind and Lindesay, 2004]. These do not avoid the ultimate problem; it can be claimed they simply postpone facing it, for one now has to ask all the same questions of origins and uniqueness about the supposed prior state to the Hot Big Bang expansion phase. The Hartle-Hawking ‘no-boundary’ proposal [Hawking, 1993] avoids the initial singularity because of a change of space-time signature, and so gets round the issue of a time of creation in an ingenious way; and Gott’s causality violation in the early universe [Gott and Li, 1998] does the same kind of thing in a different way. Such proposals cannot overcome the ultimate existential question: *Why has*

one specific state occurred rather than any of the other possibilities? How was it decided that this particular kind of universe would be the one actually instantiated? This question cannot be solved by physics alone, unless one can show that only one form of physics is self-consistent; but the variety of proposals made is evidence against that suggestion.

The explanation of initial conditions has been the aim of the family of theories one can label collectively as ‘quantum cosmology’ [Hawking, 1993; Gott and Li, 1998; Gibbons *et al.*, 2003]; however as discussed earlier, here we inevitably reach the limits to what the scientific study of the cosmos can ever say — if we assume that such studies must of necessity involve an ability to observationally or experimentally check our theories. No physical experiment at all can help here because of the uniqueness of the universe, and the feature that no spacetime exists prior to (in a causal sense) such a beginning; so brave attempts to define a ‘physics of creation’ stretch the meaning of ‘physics’. Prior to the start (if there was a start) physics as we know it is not applicable and our ordinary language fails us because time did not exist, so our natural tendency to contemplate what existed or happened ‘before the beginning’ is highly misleading — there was no ‘before’ then, indeed there was no ‘then’ then! Talking as if there was is commonplace, but quite misleading in trying to understand a scientific concept of ‘creation’ [Grunbaum, 1989]. We run full tilt into the impossibility of testing the causal mechanisms involved, when physics did not exist. No experimental test can determine the nature of any mechanisms that may be in operation in circumstances where even the concepts of cause and effect are suspect. This comes particularly to the fore in proposing ‘laws of initial conditions for the universe’ — for here we are apparently proposing a theory with only one object. Physics laws are by their nature supposed to cover more than one event, and are untestable if they do not do so (Sec. 3).

6.3 *Special or general*

The present state of the universe is very special. Explanation of the present large-scale isotropy and homogeneity of the universe means determining the dynamical evolutionary trajectories relating initial to final conditions, and then essentially either *explaining initial conditions*, where we run into difficulties (Sec. 6.2), or *showing they are irrelevant*. The issue raised is whether the universe started off in a very special geometrical state:

Thesis D3: The initial state of the universe may have been special or general. *Whether there was generality or speciality of geometrical initial conditions for the universe is a key question. It seems likely that the initial state of the observed part of the universe was not generic.*

The assumption that the universe is geometrically special was encoded in the Cosmological Principle, taken as a founding principle in cosmology until the 1960’s, i.e. as an ‘explanation’ of special initial conditions [Bondi, 1960; Weinberg, 1972]. Then Misner introduced the chaotic cosmology programme [Misner, 1968], based

on the idea of a universe with generic initial conditions being isotropised at later times by physical processes such as viscosity, making initial conditions irrelevant. This concept of isotropisation then became central to the inflationary family of theories (Sec. 2.6), with the underlying assumption being that ‘fine tuning’ of initial conditions is unphysical and to be avoided. Both programmes are however only partially successful: one can explain a considerable degree of isotropisation and homogenization of the physical universe by either process, but this will not work in all circumstances. Inflation can get rid of much anisotropy [Wald, 1983] but inhomogeneity must be restricted if inflation is to succeed in producing a universe like that we see today, and the success of inflation in solving the horizon issue for FL models — where exact homogeneity exists to start with — will not necessarily be replicated in anisotropic models. Universes that are initially too anisotropic may never inflate, and the horizon problem may not be solved in such models if they do;²⁶ and only rather special states lead to ordinary thermodynamics [Penrose, 1989a; Penrose, 2004; Wald, 2005; Carroll and Chen, 2005], which is taken to be true in inflationary physics.

Inflation can only be guaranteed to succeed if initial conditions are somewhat restricted; some degree of geometric speciality must have occurred at the start of the observed region of the universe. This special domain might possibly occur within the context of a much larger universe domain where conditions vary randomly, and only isolated regions lead to inflation and eventually domains such as that we see around us; attractive as this may be, it is an untestable hypothesis (essentially a version of the multiverse proposal, see Sec. 9.2).

Special initial conditions (which inflation proposes to explain) might have just occurred that way. The ultimate issue is that *we have no proof as to whether initial conditions for the universe were special or general; either could have occurred*. If we state these conditions must have been general, we are making a philosophical claim, for it is not a provable physical statement. Part of the problem is that we have no agreed measure on the space of possible universes; what seems special or general depends on the choice of such a measure.

7 ISSUE E: THE UNIVERSE AS THE BACKGROUND FOR EXISTENCE

The universe provides the environment for all of science, by determining the initial conditions within which all physical laws are constrained to operate, thus setting boundary conditions for all local physics. Together with suitable equations of state for the matter or structural equations for complex systems, these determine the nature of physical outcomes. The uniqueness of cosmology lies in that it considers the origin of such conditions.

²⁶Most inflationary studies show only that the *geometric* horizon problem is solved in the very special RW geometries; but there is no *physical* horizon problem in those geometries, for they are by assumption spatially homogeneous and isotropic *ab initio*.

7.1 *Laws and boundary conditions*

A fundamental assumption underlying physical cosmology is the idea that *the laws of physics are the same everywhere in the physical universe*: those we determine in a laboratory here and now will be the same as apply at very distant places (e.g. determining the astrophysics of qso's at redshift $z = 6$), at very early times (e.g. at the time of nucleosynthesis), and at very late times. Without this assumption, explanatory theories have no solid foundation. However because of the uniqueness of the universe discussed above (see Sec. 3), unlike the rest of physics where the distinction is clear and fundamental, in the cosmological context the distinction between laws and boundary conditions becomes blurred.

Thesis E1: Physical laws may depend on the nature of the universe. *We have an essential difficulty in distinguishing between laws of physics and boundary conditions in the cosmological context of the origin of the universe. Effective physical laws may depend on the boundary conditions of the universe, and may even vary in different spatial and/or temporal locations in the cosmos.*

Because we cannot vary the initial conditions in any way, as far as we are concerned they are necessary rather than contingent — so the essential distinction between initial conditions and laws is missing. The distinction is clear once the cosmos has come into existence — but we are concerned with ‘prior’ conditions associated with the creation of the cosmos and the very existence of physical laws. Certainly any proposal for distinguishing between laws of nature and boundary conditions governing solutions to those laws is untestable in this context. Given the feature that the universe is the unique background for all physics, it is therefore not far-fetched to suggest that it is possible the cosmos influences the *nature* of local physical laws, rather than just their initial conditions [Ellis and Sciama, 1972; Ellis, 2002]. This has been examined over many decades in three specific cases.

- (a) *Varying ‘constants’*: It might be that there is a time variation in physical constants of nature [Barrow, 2003] related to the expansion of the universe, as proposed in the case of the gravitational constant G by Dirac [Dirac, 1938], developed in depth by Jordan and then Brans and Dicke [Brans and Dicke, 1961]. Such proposals must be consistently developed in relation to the rest of physics and should be related to dimensionless constants, as otherwise they may simply be disguised variations in the units of measurements used, rather than being a genuine physical change (various claims that the speed of light ‘ c ’ may vary fall into this category [Ellis and Uzan, 2005]). This proposal has received impetus in recent times from ideas based in quantum field theory and string theory, suggesting that many of the ‘constants of nature’ are in fact contingent, depending on the nature of the vacuum state [Susskind, 2003; Freivogel *et al.*, 2005a]. This kind of proposal is to some degree open to observational test [Cowie and Songaila, 1995; Will, 1979], and in the cases where it has been investigated it seems that it does not occur in the visible region of the universe — the constants of nature are indeed invariant, with

one possible exception: the fine structure constant, where there is claimed to be evidence of a very small change over astronomical timescales [Barrow, 2003]. That issue is still under investigation. Testing such invariance is fundamentally important, precisely because cosmology usually assumes as a ground rule that physics is the same everywhere in the universe. If this were not true, local physics would not guide us adequately as to the behaviour of matter elsewhere or at other times, and cosmology would become an arbitrary guessing game. In order to proceed in a scientific manner when such variation is proposed, one needs then to hypothesize the manner of such variation. Thus the old laws where G was constant are replaced by new laws governing its time variation [Brans and Dicke, 1961]; the principle of nature being governed by invariant (unchanging) physical laws and associated constants remains.²⁷ Thus in the end the proposal is to replace simpler old laws by new more complex ones. These must then be assumed invariant, or we cannot proceed scientifically.

- (b) *Inertia and Mach's Principle*: It might be that the local inertial properties of matter are determined by the distant distribution of matter in the universe, so that if the universe were different, inertia would be different. This is the complex of ideas referred to as Mach's principle [Barbour and Pfister, 1995], which served as a major impetus for Einstein's cosmological ideas. The precise meaning and implications of this idea remain controversial.
- (c) *The arrow of time*: The existence and direction of the macroscopic arrow of time in physics — and hence in chemistry, biology, psychology, and society — is related to boundary conditions in the past and future of the universe. The fundamental physical laws by themselves are time symmetric, and so unable to explain this feature [Davies, 1974; Ellis and Sciama, 1972; Zeh, 1992; Uffink, 2006]. A recent argument of this kind is Penrose's claim that the existence of the arrow of time is crucially based in the universe having had rather special initial conditions [Penrose, 1989b; Penrose, 1989a; Wald, 2005]. Thus what appears in ordinary physics as an immutable law of nature (viz. the Second Law of Thermodynamics with a given arrow of time) may well be the result of specific boundary conditions at the start and end of the universe. It might not be true in all universes, even if the underlying fundamental physical laws are the same.

In each case proposals have been made as to the possible nature of the deeper underlying unchanging laws, and the relations between the state of the universe and the resultant effective laws in that context. This is also proposed in the 'landscape' of possibilities of string theory [Susskind, 2005]. These proposals are however intrinsically untestable, for the reasons explained above (Sec. 3): we cannot change

²⁷ "Despite the incessant change and dynamic of the visible world, there are aspects of the fabric of the universe which are mysterious in their unshakeable constancy. It is these mysterious unchanging things that make our universe what it is and distinguish it from other worlds we might imagine" (Barrow [Barrow, 2003], p. 3).

the boundary conditions of the universe and see what happens; but they do serve as a continuing fertile source of ideas.

7.2 *Alternative physics*

In any case, the important conclusion is that it is certainly appropriate for cosmology to consider what would have happened if, not only the boundary conditions at the beginning of the universe, but also the laws of physics had been different [Susskind, 2005]:

Thesis E2: We cannot take the nature of the laws of physics for granted.

Cosmology is interested in investigating hypothetical universes where the laws of physics are different from those that obtain in the real universe in which we live — for this may help us understand why the laws of physics are as they are (a fundamental feature of the real physical universe).

One cannot take the existence and nature of the laws of physics (and hence of chemistry) as unquestionable in cosmology — which seems to be the usual habit in biological discussions on the origin and evolution of life. This is in stark contrast to the rest of science, where we are content to take the existence and nature of the laws describing the fundamental behaviour of matter as given and unchangeable. Cosmological investigation is interested in the properties of hypothetical universes with different physical behaviour. Consideration of ‘what might have been’ is a useful cosmological speculation that may help throw light on what actually is; this is a statement of the usefulness of ‘Gedanken experiments’ in cosmology.

Indeed if one wants to investigate issues such as why life exists in the universe, consideration of this larger framework — in essence, a hypothetical ensemble of universes with many varied properties — is essential (this is of course not the same as assuming an ensemble of such universes actually exists, cf. the discussion below in Sec. 9.2). However we need to be very cautious about using any claimed statistics of universes in such a hypothetical ensemble of all possible or all conceivable universes. This is usually not well defined, and in any case is only relevant to physical processes if either the ensemble actually exists, rather than being a hypothetical one, or if it is the outcome of processes that produce well-defined probabilities — an untestable proposal. We can learn from such considerations the nature of possible alternatives, but not necessarily the probability with which they might occur (if that concept has any real meaning).

7.3 *Emergence of complexity*

As the universe evolves an increase of complexity takes place in local systems as new kinds of objects come into being that did not exist before — nuclei, atoms, stars and galaxies, planets, life, consciousness, and products of the mind such as books and computers [Morowitz, 2002]. New kinds of physical states come into being at late times such as Bose-Einstein condensates, that plausibly cannot exist without the intervention of intelligent beings.

Thesis E3: Physical novelty emerges in the expanding universe. *New kinds of physical existence come into being in the universe as it evolves, that did not exist previously. Their existence is allowed by the boundary conditions provided by the universe for local systems, together with the possibility space generated by the underlying physics. While their physical existence is novel, every new thing that comes into being is foreshadowed in possibility structures that precede their existence.*

Physical existence is new as the universe evolves, but there had to be precursors of the novel in the possibility space allowed by physics, so that they could come into being. In this sense the truly novel does not emerge *ex nihilo* but rather is discovered. The universe is the environment that allows this to happen. The nature of the features leading to the existence of life, and their possible causes, is discussed in Sec. 9.1.

8 ISSUE F: THE EXPLICIT PHILOSOPHICAL BASIS

Consequent on the discussion above, and particularly items **B6**, **C2**, and **D2**, it follows that

Thesis F1: Philosophical choices necessarily underly cosmological theory. *Unavoidable metaphysical issues inevitably arise in both observational and physical cosmology. Philosophical choices are needed in order to shape the theory.*

There is of course always a philosophical basis to any scientific analysis, namely adoption of the basic scientific method and a commitment to the attempt to explain what we see as far as possible simply in terms of causal laws, ultimately based in physics. This will clearly be true also in cosmology. However we need further explicit philosophical input in order to attain specific geometric models — for example a Copernican principle, as explained above, see Sec. 4.2.2 — and to determine what form physical cosmology should take in the very early universe, for example deciding which physical principle to use as the core of one's extrapolation of known physics to the unknown (Sec. 5). Underlying both sets of choices are criteria for satisfactoriness of a cosmological model, which help decide which feature to focus on in formulating a theory. Of particular importance is the scope chosen for our cosmological theory; together with the choice of criteria for a good theory, this is a philosophical decision that will shape the rest of the analysis. Some cosmologists tend to ignore the philosophical choices underlying their theories; but simplistic or unexamined philosophical standpoints are still philosophical standpoints!

8.1 *Criteria for theories*

As regards criteria for a good scientific theory [Kuhn, 1977], typical would be the following four areas of assessment:

1. *Satisfactory structure*: (a) internal consistency, (b) simplicity (Occam's razor), and (c) aesthetic appeal ('beauty' or 'elegance').
2. *Intrinsic explanatory power*: (a) logical tightness, (b) scope of the theory — the ability to unify otherwise separate phenomena, and (c) probability of the theory or model with respect to some well-defined measure;
3. *Extrinsic explanatory power, or relatedness*: (a) connectedness to the rest of science, (b) extendability — providing a basis for further development;
4. *Observational and experimental support*, in terms of (a) testability: the ability to make quantitative as well as qualitative predictions that can be tested; and (b) confirmation: the extent to which the theory is supported by such tests as have been made.

It is particularly the latter that characterizes a scientific theory, in contrast to other types of theories claiming to explain features of the universe and why things happen as they do. It should be noted that *these criteria are philosophical in nature in that they themselves cannot be proven to be correct by any experiment*. Rather their choice is based on past experience combined with philosophical reflection. One could attempt to formulate criteria for good criteria for scientific theories, but of course these too would need to be philosophically justified. The enterprise will end in infinite regress unless it is ended at some stage by a simple acceptance of a specific set of criteria.

Thesis F2: Criteria of satisfactoriness for theories cannot be scientifically chosen or validated. *Criteria of satisfactoriness are necessary for choosing good cosmological theories; these criteria have to be chosen on the basis of philosophical considerations. They should include criteria for satisfactory structure of the theory, intrinsic explanatory power, extrinsic explanatory power, and observational and experimental support.*

The suggestion here is that the above proposed criteria are a good set to use in investigating cosmology; they include those most typically used ([Kuhn, 1977]; and see [Penrose, 2004; Susskind, 2005] for comments on such criteria).

8.1.1 *Conflicts between criteria.*

These criteria are all acknowledged as desirable. The point then is that generally in pursuing historical sciences, and in particular in the cosmological context, they will not all be satisfied to the same degree, and may even lead to opposing conclusions:

Thesis F3: Conflicts will inevitably arise in applying criteria for satisfactory cosmological theories. *Philosophical criteria for satisfactory cosmological theories will in general come into conflict with each other, so that one will have to choose between them to some degree; this choice will shape the resulting theory.* [Ellis, 1991].

The thrust of much recent development has been away from observational tests towards strongly theoretically based proposals, indeed sometimes almost discounting observational tests. At present this is being corrected by a healthy move to detailed observational analysis of the consequences of the proposed theories, marking a maturity of the subject. However because of all the limitations in terms of observations and testing [criteria (4)], in the cosmological context we still have to rely heavily on other criteria, and some criteria that are important in most of science may not really make sense. This is true of **2(c)** in particular, as discussed above, see Sec. 3; nevertheless many approaches still give the idea of probability great weight. At a minimum, the ways this can make sense needs exploration and explication. Furthermore the meaning of some of the criteria may come into dispute. **1(b)** is clearly a case in point : for example, is the idea of an existent ensemble of universes displaying all possible behaviours simple (because it is a single idea that can be briefly stated), or immensely complex (because that statement hides all the complexities and ambiguities involved in the idea of an infinity of possibilities)? **1(c)** is also controversial('beauty is in the eye of the beholder'), see [Susskind, 2005] for a discussion.

The tenor of scientific understanding may change, altering the balance of what is considered a good explanation and what is not. An example [Ellis, 1990] is the way cosmologists strongly resisted the idea of an evolving universe in the 1920's, at a time when biological evolution was very well established but the idea of continental drift was also being strongly resisted. The change to an appreciation of the explanatory power of an evolving model came later in both cases; but even then in the cosmological case, for either aesthetic or metaphysical reasons, some still sought for a steady state description, resisting the implication of a beginning to the universe. That tendency is still with us today, in the form of models that are eternal in one way or another (e.g. some forms of chaotic inflation). Another example is the change from supposition of underlying order, expressed in the idea of a Cosmological Principle, to a broad supposition of generic disordered conditions, embodied in the ideas of inflation. Associated with this is a shift from making geometric assumptions to providing physical explanatory models. It is this shift that underlies the major present support for inflation:

Thesis F4: The physical reason for believing in inflation is its explanatory power as regards structure growth in the universe. *Inflation predicts the existence of Gaussian scale-free perturbations in the early universe thereby (given the presence of cold dark matter) explaining bottom-up structure formation in a satisfactory way. This theory has been vindicated spectacularly through observations of the CBR and matter power spectra. It is this explanatory power that makes it so acceptable to physicists, even though the underlying physics is neither well-defined nor tested, and its major large-scale observational predictions are untestable.*

The physical explanatory power of inflation in terms of structure formation, supported by the observational data on the fluctuation spectra, is spectacular. For

most physicists, this trumps the lack of identification and experimental verification of the underlying physics (Sec. 5). Inflation provides a causal model that brings a wider range of phenomena into what can be explained by cosmology (Criterion **2(b)**), rather than just assuming the initial data had a specific restricted form. Explaining flatness ($\Omega_0 \simeq 1$ as predicted by inflation) and homogeneity reinforces the case, even though these are philosophical rather than physical problems (they do not contradict any physical law; things could just have been that way). However claims on the basis of this model as to what happens very far outside the visual horizon (as in the chaotic inflationary theory) results from prioritizing theory over the possibility of observational and experimental testing [Earman and Mosterin, 1999]. It will never be possible to *prove* these claims are correct.

8.2 *The scope of cosmology*

To sensibly choose priorities for the criteria just discussed, we need an answer to the question, How much should we try to explain?

Thesis F5: Cosmological theory can have a wide or narrow scope of enquiry. *The scope we envisage for our cosmological theory shapes the questions we seek to answer. The cosmological philosophical base becomes more or less dominant in shaping our theory according to the degree that we pursue a theory with more or less ambitious explanatory aims in terms of all of physics, geometry, and underlying fundamental causation.*

This is a choice one has to make, as regards both foundations and outcomes. Given a decision on this, one can sensibly debate what is the appropriate philosophical position to adopt in studying a cosmological theory with that scope. The study of expansion of the universe and structure formation from nucleosynthesis to the present day is essential and well-informed. The philosophical stance adapted is minimal and highly plausible. The understanding of physical processes at earlier times, back to quantum gravity, is less well founded. The philosophical stance is more significant and more debatable. Developments in the quantum gravity era are highly speculative; the philosophical position adopted is dominant because experimental and observational limits on the theory are lacking.

One can choose the degree to which one will pursue the study of origins [Fabian, 1989] back to earlier and earlier times and to more fundamental causal issues, and hence the degree to which specific philosophical choices are dominant in one's theory. The basic underlying cosmological questions are [Ellis, 1991]:

1. *Why do the laws of physics have the form they do?* Issues arise such as what makes particular laws work? For example, what guarantees the behaviour of a proton, the pull of gravity? What makes one set of physical laws 'fly' rather than another? If for example one bases a theory of cosmology on string theory [Susskind, 2005], then who or what decided that quantum gravity would have a nature well described by string theory? If one considers

all possibilities, considering string theory alone amounts to a considerable restriction.

2. *Why do boundary conditions have the form they do ?* The key point here (Sec. 6.2), is how are specific contingent choices made between the various possibilities, for example whether there was an origin to the universe or not.
3. *Why do any laws of physics at all exist ?* This relates to unsolved issues concerning the nature of the laws of physics: are they descriptive or prescriptive? (Sec. 9.3.3). Is the nature of matter really mathematically based in some sense, or does it just happen that its behaviour can be described in a mathematical way?
4. *Why does anything exist ?* This profound existential question is a mystery whatever approach we take.²⁸

Finally the adventurous also include in these questions the more profound forms of the contentious Anthropic question [Carr and Rees, 1979; Davies, 1982; Barrow and Tipler, 1984; Tegmark, 1998; Susskind, 2005]:

5. *Why does the universe allow the existence of intelligent life?* This is of somewhat different character than the others and largely rests on them but is important enough to generate considerable debate in its own right.

The status of all these questions is philosophical rather than scientific, for they cannot be resolved purely scientifically. How many of them — if any — should we consider in our construction of and assessments of cosmological theories?

One option is *to decide to treat cosmology in a strictly scientific way*, excluding all the above questions, because they cannot be solved scientifically. One ends up with a solid technical subject that by definition excludes such philosophical issues. This is a consistent and logically viable option. This logically unassailable position however has little explanatory power; thus most tend to reject it because of criteria **2(b)** and **3** above.

The second option is to decide that *these questions are of such interest and importance that one will tackle some or all of them, even if that leads one outside the strictly scientific arena*. It is here that criteria **2** and **3** above are to some degree in conflict with criterion **4**. Thus if we try to explain the origin of the universe itself, these philosophical choices become dominant precisely because the experimental and observational limits on the theory are weak; this can be seen by viewing the variety of such proposals that are at present on the market.

8.3 *Limits of Representation and Knowledge of Reality*

It follows from the above discussion that there are limits to what the scientific method can achieve in explanatory terms. We need to respect these limits and

²⁸But see Grunbaum [Grunbaum, 2004] for a dissenting view.

acknowledge clearly when arguments and conclusions are based on some philosophical stance rather than purely on testable scientific argument. If we acknowledge this and make that stance explicit, then the bases for different viewpoints are clear and alternatives can be argued about rationally.

A crucial underlying feature here is relating the nature of epistemology to ontology: how do we relate evidence to our theories of existence? A further key issue is the relation of models to reality:

Thesis F6: Reality is not fully reflected in either observations or theoretical models. *Problems arise from confusion of epistemology (the theory of knowledge) with ontology (the nature of existence): existence is not always manifest clearly in the available evidence. The theories and models of reality we use as our basis for understanding are necessarily partial and incomplete reflections of the true nature of reality, helpful in many ways but also inevitably misleading in others. They should not be confused with reality itself!*

The confusion of epistemology with ontology occurs all the time, underlying for example the errors of both logical positivism and extreme relativism. In particular, it is erroneous to assume that lack of evidence for the existence of some entity is proof of its non-existence. In cosmology it is clear for example that regions may exist from which we can obtain no evidence (because of the existence of horizons); so we can sometimes reasonably deduce the existence of unseen matter or regions from a sound extrapolation of available evidence (no one believes matter ends at or just beyond the visual horizon). However one must be cautious about the other extreme, assuming existence can always be assumed because some theory says so, regardless of whether there is any evidence of existence or not. This happens in present day cosmology, for example in presentations of the case for multiverses, even though the underlying physics has not been experimentally confirmed. It may be suggested that arguments ignoring the need for experimental/observational verification of theories ultimately arise because these theories are being confused with reality, or at least are being taken as completely reliable total representations of reality. This occurs in

- Confusing computer simulations of reality with reality itself, when they can in fact represent only a highly simplified and stylized version of what actually is;
- Confusing the laws of physics themselves with their abstract mathematical representation (if indeed they are ontologically real, c.f. Sec. 10.1), or confusing a construction of the human mind ('Laws of Physics') with the reliable behaviour of ponderable matter (if they are not ontologically real);
- Confusing theoretically based outcomes of models with proven observational results (e.g. claiming the universe necessarily has flat spatial sections: $\Omega_0 = 1$, and so this can be taken for granted, when the value of Ω_0 can and should be observationally determined precisely because this then tests that prediction).

No model (literary, intuitive, or scientific) can give a perfect reflection of reality. Such models are always selective in what they represent and partial in the completeness with which they do so. The only model that would reflect reality fully is a perfect fully detailed replica of reality itself! This understanding of the limits of models and theories does not diminish the utility of these models; rather it helps us use them in the proper way. This is particularly relevant when we consider how laws of nature may relate to the origins of the universe itself, and to the existence and nature of life in the expanding universe. The tendency to rely completely on our theories, even when untested, seems sometimes to arise because we believe they are the same as reality — when at most they are *descriptions* of reality.

9 KEY ISSUES

There are some interrelated key issues where the features identified above either are at the heart of current debates, or are likely to be at the heart of future debates. They are: the reason cosmological conditions allow the existence of life (anthropic issues), the closely related issue of the possible existence of multiverses; and the natures of existence, including the questions of the existence of infinities and the nature of the laws of physics. We look at them in turn in this section. To some degree they have already been considered above, but they are specifically featured here because of the important role they will probably play in discussion in the future.

9.1 Issue G: The anthropic question: Fine tuning for life

One of the most profound fundamental issues in cosmology is the Anthropic question, see [Davies, 1982; Barrow and Tipler, 1984; Earman, 1987; Fabian, 1989; Davies, 1987; Balashov, 1991; Rees, 1999; Rees, 2003; Barrow, 2003]: *why does the Universe have the very special nature required in order that life can exist?* The point is that a great deal of “fine tuning” is required in order that life be possible. There are many relationships embedded in physical laws that are not explained by physics, but are required for life to be possible; in particular various fundamental constants are highly constrained in their values if life as we know it is to exist:

“A universe hospitable to life — what we might call a biophilic universe — has to be special in many ways ... Many recipes would lead to stillborn universes with no atoms, no chemistry, and no planets; or to universes too short lived or too empty to evolve beyond sterile uniformity” [Rees, 2003].

How has it come about that the Universe permits the evolution and existence of intelligent beings at any time or place? “What features of the universe were essential for creatures such as ourselves, and is it through coincidence or for some deeper reason that our universe has these features?” [Gribbin and Rees, 1991]. Whether one regards this as an appropriate issue for cosmology to discuss depends, as discussed above (Sec. 8.2), on the scope one envisages for cosmology. The viewpoint taken here will be that this is one of the major issues one might wish

to explain, and indeed a substantial literature considers this. Here we explore the nature of this fine tuning, and then consider possible answers as to how it arises. There are three aspects that we consider in turn (cf. [Susskind, 2005]).

9.1.1 *Laws of physics and the existence of complexity*

The laws of physics and chemistry are such as to allow the functioning of living cells, individuals, and ecosystems of incredible complexity and variety, and it is this that has made evolution possible. What requires explanation, *is why the laws of physics are such as to allow this complex functionality to work*, without which no evolution whatever would occur. We can conceive of universes where the laws of physics (and so of chemistry) were different than in ours. Almost any change in these laws will prevent life as know it from functioning.

The first requirement is *the existence of laws of physics that guarantee the kind of regularities that can underlie the existence of life*. These laws as we know them are based on variational and symmetry principles; we do not know if other kinds of laws could produce complexity. If the laws are in broad terms what we presently take them to be, the following *inter alia* need to be right, for life of the general kind we know to exist [Davies, 1982; Gribbin and Rees, 1991]:

- Quantization that stabilizes matter and allows chemistry to exist through the Pauli exclusion principle.
- The neutron-proton mass differential must be highly constrained. If the neutron mass were just a little less than it is, proton decay could have taken place so that by now no atoms would be left at all [Davies, 1982].
- Electron-proton charge equality is required to prevent massive electrostatic forces overwhelming the weaker electromagnetic forces that govern chemistry.
- The strong nuclear force must be strong enough that stable nuclei exist [Davies, 1982]; indeed complex matter exists only if the properties of the nuclear strong force lies in a tightly constrained domain relative to the electromagnetic force [Tegmark, 2003].
- The chemistry on which the human body depends involves intricate folding and bonding patterns that would be destroyed if the fine structure constant (which controls the nature of chemical bonding) were a little bit different.
- The number D of large spatial dimensions must be just 3 for complexity to exist [Tegmark, 2003; Rees, 2003].

Hogan has examined the freedom in the parameters of the standard model of particle physics and concluded that 5 of the 17 free parameters of the standard model must lie in a highly constrained domain if complex structures are to exist [Hogan, 2003]. This is of course taking the basic nature of the standard model of particle physics for granted. If this were not so, it is difficult to determine what the

constraints would be. However his study is sufficient to show that whatever the nature of fundamental physics, and in particular of particle physics, may be, only a small subset of all possible laws of physics will be compatible with the existence of complexity.

9.1.2 Laws of physics and the existence of congenial environments

The creation through astrophysical processes of suitable habitats for life to exist (the existence of planets circling stable stars, for example) depends to some degree on the nature of the fundamental physical laws. If the laws are in broad terms what we presently take them to be, the requirements for such habitats to exist include:

- The gravitational force must create large stable structures (planets and stars) that can be the habitat for life and their energy source respectively. This requires the gravitational force to be very weak relative to electrical forces. The ratio \mathcal{N} of the strength of the electromagnetic force to the gravitational force must be close to the observed value: $\mathcal{N} \simeq 10^{36}$ [Rees, 1999, Ch. 3].
- The weak force must allow helium production that leaves sufficient hydrogen over; it is related to gravity through a numerical factor of 10^{-11} , which cannot be much different. And for this to work, the neutron-proton mass difference must be close to the mass of the electron [Davies, 1982].
- A stellar balance should allow a long lifetime for stars like the sun, so allowing the transmutation of the light elements into heavy elements. This requires that the nuclear fusion efficiency \mathcal{E} be close to the observed value: $\mathcal{E} \simeq 0.007$ [Rees, 1999, Ch. 4].
- One needs to overcome the beryllium “bottleneck” in the making of heavy elements through nuclear reactions in stars [Gribbin and Rees, 1991; Susskind, 2005]. The production of carbon and oxygen in stars requires the careful setting of two different nuclear energy levels to provide a resonance; if these levels were just a little different, the elements we need for life would not exist [Fabian, 1989]. Indeed it was on this basis that Hoyle famously predicted a carbon-12 energy level that has since been experimentally confirmed.
- One needs something like the existence of neutrinos and the weak interaction with its specific coupling constant in order to underly supernovae explosions that spread heavy elements through space, as seeds for planetary formation [Gribbin and Rees, 1991].
- The nuclear force must be weak enough that di-protons do not exist, otherwise no protons will be left over to enable heavier elements to exist [Davies, 1982].
- The neutrino mass must not be too high, or the universe will not last long enough [Davies, 1982].

9.1.3 Cosmological boundary/initial conditions and congenial environments

Finally, given laws of physics that are suitable in terms of satisfying the requirements of both the previous sections, the universe itself must also be suitable, in terms of its initial or boundary conditions, for life to exist. If the laws of physics are basically the same as we now believe them to be, these cosmological requirements include

- The size of the universe and its age must be large enough. There could be universes that expanded and then recollapsed with a total lifetime of only 100,000 years; we need a sufficiently old universe for second generation stars to come into existence and then for planets to have a stable life for long enough that evolution could lead to the emergence of intelligent life. Thus the universe must be at about 15 billion years old for life to exist [Gribbin and Rees, 1991], hence we must have $\Omega_{matter} \simeq 0.3$ [Rees, 1999, Ch. 6].
- The size of the cosmological constant must not be too large, or galaxies will not form; we need $|\Omega_\Lambda| < 1$ for galaxies to exist [Rees, 1999, Ch. 7]; [Susskind, 2005].
- The seeds in the early universe for fluctuations that will later grow into galaxies must be of the right size that structures form without collapsing into black holes: the number Q characterizing the size of primordial ripples on the LSS (and hence the geometry of the perturbed cosmological model, see Sec. 2.5.2) must therefore be of the order $Q \simeq 10^{-5}$ [Rees, 1999, Ch. 8].

The complex of interacting systems in a human body could not possibly work if a series of delicate conditions were not maintained. For example, the background radiation might never drop below 3000 K, so that matter was always ionized (electrons and nuclei always remaining separate from each other); the molecules of life could then never form. Black holes might be so common that they rapidly attracted all the matter in the universe, and there never was a stable environment in which life could develop. Cosmic rays could always be so abundant that any tentative organic structures were destroyed before they could replicate. Overall,

- There must be non-interference with local systems. The concept of locality is fundamental, allowing local systems to function effectively independently of the detailed structure of the rest of the Universe. We need the universe and the galaxies in it to be largely empty, and gravitational waves and tidal forces to be weak enough,²⁹ so that local systems can function in a largely isolated way [Ellis, 2002].
- The fact that the night sky is dark ('Olbers' paradox' [Bondi, 1960; Harrison, 2000]) is a consequence of the expansion of the universe together with the photon to baryon ratio. This feature is a necessary condition for the existence

²⁹Thus the Weyl tensor C_{abcd} must be suitably small everywhere, presumably implying an almost-RW geometry, cf. [Stoeger *et al.*, 1995].

of life: the biosphere on Earth functions by disposing of waste energy to the heat sink of the dark night sky [Penrose, 1989b]. Thus one way of explaining why the sky is observed to be dark at night is that if this were not so, we would not be here to observe it.

- The existence of the arrow of time, and hence of laws like the second law of thermodynamics, are probably necessary for evolution and for consciousness. This depends on boundary conditions at the beginning and end of the Universe (Sec. 7.1).
- Presumably the emergence of a classical era out of a quantum state is required. The very early universe would be a domain where quantum physics would dominate, leading to complete uncertainty and an inability to predict the consequence of any initial situation; we need this to evolve to a state where classical physics leads to the properties of regularity and predictability that allow order to emerge.
- Physical conditions on planets must be in a quasi-equilibrium state for long enough to allow the delicate balances that enable our existence, through the very slow process of evolution, to be fulfilled.

Thus the existence of suitable local systems to be a habitat for life depends critically on the large-scale properties of very distant matter. These provides a stable local environment within which life can develop.

9.1.4 *Fine tuning overall*

Thus there are many ways that conditions in a universe could prevent life occurring. Life will occur only if: there exist heavy elements; there is sufficient time for evolution of advanced life forms to take place; there are regions in the universe that are neither too hot nor too cold; there are precisely restricted values of the fundamental constants that control chemistry and local physics; and so on. These conditions will not be true in a generic universe. In summary,

Thesis G1: Life is possible because both the laws of physics and the boundary conditions for the universe have a very special nature. *Only particular laws of physics, and particular initial conditions in the Universe, allow the existence of intelligent life of the kind we know. No evolutionary process whatever is possible for any kind of life if these laws and conditions do not have this restricted form.*

Why is this so? One should note that we can only meaningfully refer here to ‘life as we know it’. One of the recurring issues is whether there could be some other quite different basis for life. You can if you wish speculate that life might exist in some immaterial form, or based only on light elements, or existent deep in empty space without the need for stars or planets to provide a viable habitat. The anthropic literature is based on assuming this is not viable, but we cannot *prove* anything in this regard. We have no idea of any basis by which life might

come into existence other than the broad principles we see in the life around us. The basic principles of life as we understand it require a great degree of complex organization enabling it to fulfil a complex variety of functions that can only, as far as we know, be based in material existence with information storage, energy usage, sensing of the external world, etc., which requires at a minimum heavy elements (carbon, nitrogen, oxygen, phosphorus for example), a long-term energy source (such as the flow of energy from the sun), and a stable environment (such as the surface of a planet). When we abandon this basis for understanding — saying ‘yes but some other form of life might exist’ without providing any proposal for its possible structure — one enters the unprofitable realm of speculation. It does not seem to provide any useful way forward.

9.1.4.1 The Weak Anthropic Principle. There are two purely scientific approaches to the Anthropic issue.³⁰ The first is the *Weak Anthropic Principle* (WAP), based on the comment: it is not surprising the observed Universe admits the existence of life, for the Universe cannot be observed unless there are observers in it [Barrow and Tipler, 1984; Balashov, 1991]. This seemingly empty statement gains content when we turn it round and ask, at what times and places in the Universe can life exist, and what are the inter-connections that are critical for its existence? It could not for example exist too early in the present expansion phase, for the night sky would then have been too hot. Furthermore one can deduce various necessary relations between fundamental quantities in order that the observers should exist (e.g. those mentioned above), so that if for example the fundamental constants vary with time or place in the Universe, life will only be possible in restricted regions where they take appropriate Anthropic values.

Hence this view basically interprets the Anthropic principle as a selection principle: the necessary conditions for observers to exist restricts the times and places from which the Universe can be observed. Because it is quite possible that conditions would not be right for life to exist anywhere in an arbitrarily selected universe, it is also usually conjoined with the idea of the existence of a multiverse, as discussed below, see Sec. 9.2. This is an interesting and often illuminating viewpoint. For example, neither the Chaotic Inflationary Universe idea (Sec. 2.6) nor any other multiverse proposal works unless we add such an Anthropic component into their interpretation to explain why we observe the Universe from a viewpoint where life exists. It is now used by some physicists to explain the low value of the cosmological constant (which quantum field theory predicts should have a very much larger value than observed, see Sec. 9.2.5), and occurs in the context of the possibility landscape of string theory [Susskind, 2005].

³⁰I omit the so-called *Final Anthropic Principle* (FAP for short), which maintains that intelligent life must necessarily evolve and then remain in existence until the end of the universe, for I do not believe it merits serious discussion as a scientific proposal; indeed it led to a famous book review referring to the *Completely Ridiculous Anthropic Principle* (CRAP for short) [Gardner, 1986].

9.1.4.2 The Strong Anthropic Principle. By contrast, the *Strong Anthropic Principle* (SAP) [Barrow and Tipler, 1984; Balashov, 1991] claims that it is necessary that intelligent life exist in the Universe; the presence of life is required in order that a universe model make sense. This is clearly a very controversial claim, for it is hard to provide scientific reasons to support this view. One can suggest that the most solid justification attempted is through the claim that the existence of an observer is necessary in order that quantum theory can make sense. However, this justification is based on one of a number of different interpretations of quantum theory; the nature of these quantum foundations is controversial, and not resolved [Isham, 1997; Dickson, 2006; Landsman, 2006].

Furthermore if we were to suppose this justification correct, then the next step is to ask: Why does the Universe need quantum mechanics anyway? The argument would be complete only if we could prove that quantum mechanics was absolutely necessary for every self-consistent Universe; but that line of reasoning cannot be completed at present, not least because quantum mechanics itself is not a fully self-consistent theory. Apart from the conceptual problems at its foundation due to the unresolved measurement issue [Isham, 1997], it suffers from divergences that so far have proved irremediable in the sense that we can work our way round them to calculate what we need, but cannot remove them. The SAP proposal has no accepted physical foundation, and also raises problematic philosophical issues [Earman, 1987]. I will not pursue it further here.

9.1.5 *The relation to fundamental physical theories*

Many physicists go further, rejecting any Anthropic form of reasoning. They regard it as a cop-out resorted to when physical theories fail to give the needed answers, and seek to obtain a full answer from physics alone [Scott, 2005; Susskind, 2005]. One possibility is that there is a fundamental theory of everything that determines the nature of physics completely, with no arbitrary parameters left, and this still to be discovered theory just happens to be of such a nature as to admit life.

However in this case the Anthropic issue returns with a vengeance: *How could it be that such a theory, based for example on variational principles and the specific invariance groups of particle physics, could just happen to lead to biophilic parameter values?* There is no clear way to answer such a question. Uniqueness of fundamental physics resolves the parameter freedom only at the expense of creating an even deeper mystery, with no way of resolution apparent. In effect, the nature of the unified fundamental force would be pre-ordained to allow, or even encourage, the existence of life; but there would be no apparent reason why this should be so.

A second possibility is that physics allows many effective theories with varying parameters — some form of multiverse, as for example may be implied by string theory [Susskind, 2003; Freivogel *et al.*, 2005a; Susskind, 2005]. If these varying options are all equally real, life can occur because in some cases the parameters will

lie in the restricted biophilic regime. Thus from this viewpoint the Anthropic idea is intimately linked with the existence of multiverses, which provide a legitimate domain for their application. We will turn to an examination of multiverses in the next section, but before doing so we will consider the range of metaphysical options for resolving the anthropic question.

9.1.6 *The metaphysical options*

To make progress on the Anthropic issue, we have to seriously consider the nature of ultimate causation: What is the fundamental cause for the phenomena we see? If we pursue the chain of physical cause and effect to its conclusion, we are still left with the question: *Why did this occur, and not something else?* Whatever the reason is, it is the ultimate cause we are seeking. Note that we are here leaving the terrain of science itself, and starting to probe the domain of metaphysics — the foundations of science and indeed of existence. As noted above, one can simply decide not to pursue such issues. If we do continue to question, there appear to be basically six approaches to the issue of ultimate causation: namely Random Chance, Necessity, High Probability, Universality, Cosmological Natural Selection, and Design. We briefly consider these in turn.

Option 1: *Random Chance, signifying nothing.* The initial conditions in the Universe just happened, and led to things being the way they are now, by pure chance. Probability does not apply. There is no further level of explanation that applies; searching for ‘ultimate causes’ has no meaning.

This is certainly logically possible, but not satisfying as an explanation, as we obtain no unification of ideas or predictive power from this approach. Nevertheless some implicitly or explicitly hold this view.

Option 2: *Necessity.* Things have to be the way they are; there is no other option. The features we see and the laws underlying them are demanded by the unity of the Universe: coherence and consistency require that things must be the way they are; the apparent alternatives are illusory. Only one kind of physics is self-consistent: all logically possible universes must obey the same physics.

To really prove this would be a very powerful argument, potentially leading to a self-consistent and complete scientific view. But we can imagine alternative universes! — why are they excluded? Furthermore we run here into the problem that we have not succeeded in devising a fully self-consistent view of physics: neither the foundations of quantum physics nor of mathematics are on a really solid consistent basis. Until these issues are resolved, this line cannot be pursued to a successful conclusion.

Option 3: *High probability.* Although the structure of the Universe appears very improbable, for physical reasons it is in fact highly probable.

These arguments are only partially successful, even in their own terms. They run into problems if we consider the full set of possibilities: discussions proposing this

kind of view actually implicitly or explicitly restrict the considered possibilities *a priori*, for otherwise it is not very likely the Universe will be as we see it. Besides, we do not have a proper measure to apply to the set of initial conditions, enabling us to assess these probabilities. Furthermore, as discussed above, see Sec. 3, application of probability arguments to the Universe itself is dubious, because the Universe is unique. Despite these problems, this approach has considerable support in the scientific community, for example it underlies the chaotic inflationary proposal (Sec. 2.6). It attains its greatest power in the context of the assumption of universality:

Option 4: *Universality*. This is the stand that “All that is possible, happens”: an ensemble of universes or of disjoint expanding universe domains is realized in reality, in which all possibilities occur [Rees, 1999; Rees, 2003; Tegmark, 2003]. In its full version, the anthropic principle is realized in both its strong form (if all that is possible happens, then life must happen) and its weak form (life will only occur in some of the possibilities that are realized; these are picked out from the others by the WAP, viewed as a selection principle). There are four ways this has been pursued.

1. *Spatial variation.* The variety of expanding universe domains is realised in space through random initial conditions, as in chaotic inflation (Sec. 2.6). While this provides a legitimate framework for application of probability, from the viewpoint of ultimate explanation it does not really succeed, for there is still then one unique Universe whose (random) initial conditions need explanation. Initial conditions might be globally statistically homogeneous, but also there could be global gradients in some physical quantities so that the Universe is not statistically homogeneous; and these conditions might be restricted to some domain that does not allow life. It is a partial implementation of the ensemble idea; insofar as it works, it is really a variant of the “high probability” idea mentioned above. If it was the more or less unique outcome of proven physics, then that would provide a good justification; but the physics underlying such proposals is not even uniquely defined, much less tested. Simply claiming a particular scalar field with some specific stated potential exists does not prove that it exists!
2. *Time variation.* The variety of expanding universe domains could be realised across time, in a universe that has many expansion phases (a Phoenix universe), whether this occurs globally or locally. Much the same comments apply as in the previous case.
3. *Quantum Mechanical.* It could occur through the existence of the Everett-Wheeler “many worlds” of quantum cosmology, where all possibilities occur through quantum branching [Deutsch, 1998]. This is one of the few genuine alternatives proposed to the Copenhagen interpretation of quantum mechanics, which leads to the necessity of an observer, and so potentially to the

Strong Anthropic interpretation considered above (see Sec. 9.1). The many-worlds proposal is controversial: it occurs in a variety of competing formulations [Isham, 1997], none of which has attained universal acceptance. The proposal does not provide a causal explanation for the particular events that actually occur: if we hold to it, we then have to still explain the properties of the particular history we observe (for example, why does our macroscopic universe have high symmetries when almost all the branchings will not?). And above all it is apparently untestable: there is no way to experimentally prove the existence of all those other branching universes, precisely because the theory gives the same observable predictions as the standard theory.

4. *Completely disconnected.* They could occur as completely disconnected universes: there really is an ensemble of universes in which all possibilities occur, without any connection with each other [Lewis, 1986; Rees, 2003; Tegmark, 2003]. A problem that arises then is, What determines what is possible? For example, what about the laws of logic themselves? Are they inviolable in considering all possibilities? We cannot answer, for we have no access to this multitude of postulated worlds. We explore this further below (Sec. 9.2).

In all these cases, major problems arise in relating this view to testability and so we have to query the meaningfulness of the proposals as scientific explanations. They all contradict Occam's razor: we "solve" one issue at the expense of envisaging an enormously more complex existential reality. Furthermore, they do not solve the ultimate question: *Why does this ensemble of universes exist?* One might suggest that ultimate explanation of such a reality is even more problematic than in the case of single universe. Nevertheless this approach has an internal logic of its own which some find compelling. We consider this approach further below, see Sec. 9.2.

Option 5: *Cosmological Natural Selection.* If a process of re-expansion after collapse to a black hole were properly established, it opens the way to the concept not merely of evolution of the Universe in the sense that its structure and contents develop in time, but in the sense that the Darwinian selection of expanding universe regions could take place, as proposed by Smolin [Smolin, 1992]. The idea is that there could be collapse to black holes followed by re-expansion, but with an alteration of the constants of physics through each transition, so that each time there is an expansion phase, the action of physics is a bit different. The crucial point then is that some values of the constants will lead to production of more black holes, while some will result in less. This allows for evolutionary selection favouring the expanding universe regions that produce more black holes (because of the favourable values of physical constants operative in those regions), for they will have more "daughter" expanding universe regions. Thus one can envisage natural selection favouring those physical constants that produce the maximum number of black holes.

The problem here is twofold. First, the supposed ‘bounce’ mechanism has never been fully explicated. Second, it is not clear — assuming this proposed process can be explicated in detail — that the physics which maximizes black hole production is necessarily also the physics that favours the existence of life. If this argument could be made water-tight, this would become probably the most powerful of the multiverse proposals.

Option 6: *Purpose or Design.* The symmetries and delicate balances we observe require an extraordinary coherence of conditions and cooperation of causes and effects, suggesting that in some sense they have been purposefully designed. That is, they give evidence of intention, both in the setting of the laws of physics and in the choice of boundary conditions for the Universe. This is the sort of view that underlies Judaeo-Christian theology. Unlike all the others, it introduces an element of meaning, of signifying something. In all the other options, life exists by accident; as a chance by-product of processes blindly at work.

The prime disadvantage of this view, from the scientific viewpoint, is its lack of testable scientific consequences (“Because God exists, I predict that the density of matter in the Universe should be x and the fine structure constant should be y ”). This is one of the reasons scientists generally try to avoid this approach. There will be some who will reject this possibility out of hand, as meaningless or as unworthy of consideration. However it is certainly logically possible. The modern version, consistent with all the scientific discussion preceding, would see some kind of purpose underlying the existence and specific nature of the laws of physics and the boundary conditions for the Universe, in such a way that life (and eventually humanity) would then come into existence through the operation of those laws, then leading to the development of specific classes of animals through the process of evolution as evidenced in the historical record. Given an acceptance of evolutionary development, it is precisely in the choice and implementation of particular physical laws and initial conditions, allowing such development, that the profound creative activity takes place; and this is where one might conceive of design taking place.³¹

However from the viewpoint of the physical sciences *per se*, there is no reason to accept this argument. Indeed from this viewpoint there is really no difference between design and chance, for they have not been shown to lead to different physical predictions.

9.1.7 *Metaphysical Uncertainty*

In considering ultimate causation underlying the anthropic question, in the end we are faced with a choice between one of the options above. As pointed out already by Kant and Hume, although we may be able to argue strongly for one or other of them, we cannot *prove* any of the options are correct [Hume, 1993].

³¹This is not the same as the view proposed by the ‘Intelligent Design’ movement. It does not propose that God tweaks the outcome of evolutionary processes.

Thesis G2: Metaphysical uncertainty remains about ultimate causation in cosmology. *We cannot attain certainty on the underlying metaphysical cosmological issues through either science or philosophy.*

If we look at the anthropic question from a purely scientific basis, we end up without any resolution, basically because science attains reasonable certainty by limiting its considerations to restricted aspects of reality; even if it occasionally strays into the area of ultimate causation, it is not designed to deal with it. By itself, it cannot make a choice between these options; there is no relevant experiment or set of observations that can conclusively solve the issue. Thus a broader viewpoint is required to make progress, taking into account both the scientific and broader considerations. The issue is of a philosophical rather than scientific nature. One important issue that then arises is what kind of data is relevant to these philosophical choices, in addition to that which can be characterized as purely scientific data (Sec. 9.3.4).

9.2 Issue H: The possible existence of multiverses

If there is a large enough ensemble of numerous universes with varying properties, it may be claimed that it becomes virtually certain that some of them will just happen to get things right, so that life can exist; and this can help explain the fine-tuned nature of many parameters whose values are otherwise unconstrained by physics [Rees, 1999; Rees, 2003]. As discussed in the previous section, there are a number of ways in which, theoretically, multiverses could be realized [Lewis, 1986; Tegmark, 2003]. They provide a way of applying probability to the universe [Sciama, 1971; Bostrom, 2002] (because they deny the uniqueness of the universe). However, there are number of problems with this concept. Besides, this proposal is observationally and experimentally untestable; thus its scientific status is debatable.

9.2.1 Definition

In justifying multiverses, it is often stated that ‘all that can occur, occurs’ (or similarly). However that statement does not adequately specify a multiverse. To define a multiverse properly requires two steps [Ellis *et al.*, 2004]. First, one needs to specify what is conceived of in the multiverse, by defining a *possibility space*: a space \mathcal{M} of all possible universes, each of which can be described in terms of a set of states s in a state space \mathcal{S} . Each universe m in \mathcal{M} will be characterized by a set of distinguishing parameters p , which are coordinates on \mathcal{S} . Choices are needed here. In geometrical terms, will it include only Robertson–Walker models, or more general ones (e.g. Bianchi models, or models without symmetries)? In gravitational terms, will it include only General Relativity, or also brane theories, models with varying G , loop quantum gravity models, string theory models with their associated possibility ‘landscapes’, and models based on the wave function of the universe concept? Will it allow only standard physics but

with varying constants, or a much wider spectrum of physical possibilities, e.g. universes without quantum theory, some with five fundamental forces instead of four, and others with Newtonian gravity? Defining the possibility space means making some kind of assumptions about physics and geometry that will then apply across the whole family of models considered possible in the multiverse, and excluding all other possibilities.

Second, one needs to specify which of the possible universes are physically realized in the multiverse, and how many times each one occurs. *A multiverse must be a physically realized multiverse and not a hypothetical or conceptual one if it is to have genuine explanatory power.* Thus one needs a distribution function $f(m)$ specifying how many times each type of possible universe m in \mathcal{M} is realised. The function $f(m)$ expresses the contingency in any actualization. Things could have been different! Thus, $f(m)$ describes a specific *ensemble of universes* or *multiverse* envisaged as being realised out of the set of possibilities. For example, $f(m)$ might be non-zero for all possible values of all the parameters p ('all that can happen, happens'); but it could be that f describes a multiverse where there are 10^{100} identical copies of one particular universe (the realization process finds a particularly successful recipe, and then endlessly replicates it).

Additionally we need a measure $d\pi$ that enables this function to determine numbers and probabilities of various properties in the multiverse: the number of universes corresponding to a set of parameter increments will be dN given by

$$(38) \quad dN = f(m)d\pi$$

for continuous parameters; for discrete parameters, we add in the contribution from all allowed parameter values. The total number of universes N in the ensemble will be given by

$$(39) \quad N = \int_{\mathcal{M}} f(m)d\pi$$

(which will often diverge), where the integral ranges over all allowed values of the member parameters and we take it to include all relevant discrete summations. The expectation value P of a quantity $p(m)$ defined on the set of universes will be given by

$$(40) \quad P = \int_{\mathcal{M}} p(m)f(m)d\pi.$$

These three elements (the possibility space, the measure, and the distribution function) must all be clearly defined in order to give a proper specification of a multiverse [Ellis *et al.*, 2004]. This is almost never done.

9.2.2 Non-uniqueness: Possibilities

There is non-uniqueness at both steps. Stating "all that is possible, happens" does not resolve what is possible. The concept of multiverses is not well defined until the

space of possible universes has been fully characterized; it is quite unclear how to do this uniquely. The issue of what is to be regarded as an ensemble of ‘all possible’ universes can be manipulated to produce any result you want, by redefining what is meant by this phrase — standard physics and logic have no necessary sway over them: what I envisage as ‘possible’ in such an ensemble may be denied by you. What super-ordinate principles are in operation to control the possibilities in the multiverse, and why? A key point here is that *our understandings of the possibilities are always of necessity arrived at by extrapolation from what we know*, and my imagination may be more fertile than yours, and neither need correspond to what really exists out there — if indeed there is anything there at all. Do we include only

- *Weak variation*: e.g. only the values of the constants of physics are allowed to vary? This is an interesting exercise but is certainly not an implementation of the idea ‘all that can happen, happens’. It is an extremely constrained set of variations.
- *Moderate variation*: different symmetry groups, or numbers of dimensions, etc. We might for example consider the possibility landscapes of string theory [Freivogel *et al.*, 2005b] as realistic indications of what may rule multiverses [Susskind, 2003; Freivogel *et al.*, 2005a; Susskind, 2005]. But that is very far indeed from ‘all that is possible’, for that should certainly include spacetimes not ruled by string theory.
- *Strong variation*: different numbers and kinds of forces, universes without quantum theory or in which relativity is untrue (e.g. there is an aether), some in which string theory is a good theory for quantum gravity and others where it is not, some with quite different bases for the laws of physics (e.g. no variational principles).
- *Extreme variation*: universes where physics is not well described by mathematics; with different logic; universes ruled by local deities; allowing magic as in the Harry Potter series of books; with no laws of physics at all? Without even mathematics or logic?

Which is claimed to be the properties of the multiverse, and why? We can express our dilemma here through the paradoxical question: *Are the laws of logic necessary in all possible universes?*

9.2.3 *Non-uniqueness: existence and causation*

A specific multiverse is defined by specifying the distribution function $f(m)$ of actually realized universes. It is unclear what mechanism can underlie such a distribution, and any proposal for such a mechanism is completely untestable. We need some indication as to *what determines existence within the possibilities defined by the supposed possibility space*: What decides how many times each one happens?

Unless we understood the supposed underlying mechanisms we can give no serious answer; and there is no prospect whatever of testing any proposed mechanism. The mechanisms supposed to underlie whatever regularities there are in the multiverse must pre-exist the existence of not merely this universe but also every other one. If one assumes a universe that is connected in the large but is locally separated into causally disconnected domains with different physical properties (as in chaotic inflation), one attains a plausible picture of a creation mechanism that can underlie an effective multiverse — but at the expense of supposing the validity of untested and perhaps untestable physics. Because of this one does not obtain a specification of a unique multiverse: the physics could be different than what we assumed.

9.2.4 *Explanatory power*

What explanatory power do we get in return for these problems? It has been suggested they explain the parameters of physics and of cosmology and in particular the very problematic observed value of the cosmological constant [Weinberg, 2000a; Weinberg, 2000b; Susskind, 2005]. The argument goes as follows: assume a multiverse exists; observers can only exist in one of the highly improbable biophilic outliers where the value of the cosmological constant is very small [Hartle, 2004]. A similar argument has been proposed for neutrino masses [Tegmark *et al.*, 2003]. If the multiverse has many varied locations with differing properties, that may indeed help us understand the Anthropic issue: some regions will allow life to exist, others will not [Barrow and Tipler, 1984; Leslie, 1989]. This does provide a useful modicum of explanatory power. However it is far from conclusive. Firstly, it is unclear why the multiverse should have the restricted kinds of variations of the cosmological constant assumed in the various analyses mentioned. If we assume ‘all that can happen, happens’ the variations will not be of that restricted kind; those analyses will not apply.

Secondly, ultimate issues remain: Why does this unique larger whole have the properties it does? *Why this multiverse rather than any other one?* Why is it a multiverse that allows life to exist? Many multiverses will not allow any life at all. To solve this, we can propose an *ensemble of ensembles of universes*, with even greater explanatory power and even less prospect of observational verification; and so on. The prospect of an infinite regress looms. Indeed if we declare (as suggested at the start of this article) that ‘the Universe’ is the total of all that physically exists, then when an ensemble of expanding universe domains exists, whether causally connected or not, that ensemble itself should be called ‘the Universe’, for it is then the totality of physically existing entities. All the foundational problems for a single existing universe domain recur for the multiverse — because when properly considered, it is indeed the Universe!

9.2.5 *Testability*

If an ensemble exists with members not connected in any physical way to the observable universe, then we cannot interact with them in any way nor observe

them, so we can say anything we like about them without fear of disproof.³² Thus any statements we make about them can have no solid scientific or explanatory status; they are totally vulnerable to anyone else who claims an ensemble with different properties (for example claiming different kinds of underlying logics are possible in their ensemble, or claiming many physically effective gods and devils in many universes in their ensemble).

Thesis H1: Multiverse proposals are unprovable by observation or experiment, but some self-consistency tests are possible. *Direct observations cannot prove or disprove that a multiverse exists, for the necessary causal relations allowing observation or testing of their existence are absent. Their existence cannot be predicted from known physics, because the supposed causal or pre-causal processes are either unproven or indeed untestable. However some self-consistency conditions for specific multiverse models can be tested.*

Any proposed physics underlying a multiverse proposal, such as Coleman-de Luccia tunneling [Coleman and de Luccia, 1980], will be an extrapolation of known physics; but the validity of that major extrapolation to cosmology is untestable.

Attempts have been made to justify the existence of multiverses as testable firstly via Rees' 'slippery slope' argument [Rees, 2003]. This runs as follows: we can reasonably assume galaxies that we cannot see exist outside the visual horizon (Sec. 8.3); why not extend this argument by small steps to show totally disconnected universes exist? The problem is that this assumes a continuity of existence that does not hold good. The domain outside our horizon is assumed to exist with similar properties to those inside because they are a continuous extension of it and have a largely common causal origin; their nature can be inferred from what we can see. Disconnected multiverse domains are assumed to have quite different properties, and their nature cannot be inferred from what we can see as there is no continuity or causal connection.

Secondly, several authors (Leslie [Leslie, 1989], Weinberg [Weinberg, 2000a; Weinberg, 2000b], and Rees [Rees, 2003] for example) have used arguments based on the idea that the universe is no *more* special than it has to be; a form of "speciality argument." According to Rees, if our universe turns out to be *even more specially* tuned than our presence requires, the existence of a multiverse to explain such "over-tuning" would be refuted; but the actual universe is not more special than this, so the multiverse is not refuted.

In more detail: naive quantum physics predicts the cosmological constant Λ to be very large. But our presence in the universe requires it to be small enough that galaxies and stars can form, so Λ must obviously be below that galaxy-forming threshold. If our universe belongs to an ensemble in which Λ was equally likely to take any value in the biophilic region (the uniform probability assumption),³³ then

³²But there are counter arguments by Leibniz [Wilson, 1989] and Lewis [Lewis, 1986, section 2.4, pp. 108–115].

³³The probability distribution for Λ will plausibly peak far away from the biophilic region, tailing down to a low value that will be approximately constant in that narrow region, cf. [Hartle, 2004].

we would not expect it to be too far below this threshold. This is because, if it's too far below the threshold, the probability of randomly choosing that universe in the ensemble becomes very small — there are very few universes with such small values of Λ in the biophilic subset of the ensemble. That is, it would be more likely that any bio-friendly universe in the ensemble would have a value of Λ closer to the threshold value. Present data on this value indicates that it is not too far below the threshold. Thus, our universe is not markedly more special that it needs to be as far as Λ is concerned, and so explaining its fine-tuning by existence of a multiverse is legitimate.

Is this argument compelling? It is a reasonable test of consistency for a multiverse that is known to exist, so that probability considerations apply; but they do not apply if there is no multiverse (Sec. 3). Additionally, probability considerations cannot ever be *conclusive*. Indeed,

Thesis H2: Probability-based arguments cannot demonstrate the existence of multiverses. *Probability arguments cannot be used to prove the existence of a multiverse, for they are only applicable if a multiverse exists. Furthermore probability arguments can never prove anything for certain, as it is not possible to violate any probability predictions, and this is a fortiori so when there is only one case to consider, so that no statistical observations are possible.*

All one can say on the basis of probability arguments is that some specific state is very improbable. But this does not prove it is impossible, indeed if it is stated to have a low probability, that is precisely a statement that it is possible. Thus such arguments can at best only give plausibility indications even when they are applicable. The assumption that probability arguments can be conclusive is equivalent to the claim that the universe is generic rather than special; but whether this is so or not is precisely the issue under debate (see Thesis **D3**). The argument is useful as a plausibility argument for a multiverse, but is not *proof* of its existence.

Finally, it has been proposed that the existence of multiverses is an inevitable consequence of the universe having infinite space sections [Tegmark, 2003; Seife, 2004], because that leads to infinite spatial repetition of conditions (cf. [Ellis and Brundrit, 1979]). But this supposed spatial infinity is an untested philosophical assumption, which certainly cannot be observationally proven to be correct. Apart from the existence of horizons preventing confirmation of this supposition, even if the entire universe were observable, proving it correct would still not be possible because by definition counting an infinite number of objects takes an infinite amount of time. This is an untestable philosophical argument, not an empirically testable one; furthermore, it can be argued to be implausible (Sec. 9.3.2). Indeed current data suggest it is not the case; this is the one good consistency test one can use for some multiverse proposals (Sec. 9.2.7).

9.2.6 Explanation vs Testability

The argument that this infinite ensemble actually exists can be claimed to have a certain explanatory economy, although others would claim that Occam's razor has

been completely abandoned in favour of a profligate excess of existential multiplicity, extravagantly hypothesized in order to explain the one universe that we do know exists. Certainly the price is a lack of testability through either observations or experiment — which is usually taken to be an essential element of any serious scientific theory.³⁴ It is not uniquely definable nor determinable, and there is a complete loss of verifiability. There is no way to determine the properties of any other universe in the multiverse if they do indeed exist, for they are forever outside observational reach. The point is that there is not just an issue of showing a multiverse exists. If this is a scientific proposition one needs to be able to show which specific multiverse exists; but there is no observational way to do this. Indeed if you can't show *which particular* one exists, it is doubtful you have shown *any* one exists.

What does a claim for such existence mean in this context? Gardner puts it this way: “There is not the slightest shred of reliable evidence that there is any universe other than the one we are in. No multiverse theory has so far provided a prediction that can be tested. As far as we can tell, universes are not even as plentiful as even *two* blackberries” [Gardner, 2003].³⁵

Thesis H3: Multiverses are a philosophical rather than scientific proposal. *The idea of a multiverse provides a possible route for the explanation of fine tuning. But it is not uniquely defined, is not scientifically testable apart from some possible consistency tests, and in the end simply postpones the ultimate metaphysical questions.*

The definitive consistency tests on some multiverse proposals (Sec. 9.2.7) are *necessary* conditions for those specific multiverse proposals, but are hardly by themselves indications that the multiverse proposal is true. The drive to believe this is the case comes from theoretical and philosophical considerations (see e.g. [Susskind, 2005]) rather than from data. The claim an ensemble physically exists³⁶ is problematic as a proposal for scientific explanation, if science is taken to involve testability. Indeed, adopting these explanations is a triumph of theory over testability [Gardner, 2003], but the theories being assumed are not testable. It is therefore a metaphysical choice made for philosophical reasons. That does not mean it is unreasonable (it can be supported by quite persuasive plausibility arguments); but its lack of scientific status should be made clear.

³⁴In [Stoeger *et al.*, 2004], the framework and conditions under which the multiverse hypothesis would be testable within a retroductive framework, given the rigorous conditions formulated in that paper; are indicated; these conditions are not fulfilled.

³⁵This contrasts strongly, for example, with Deutsch's and Lewis's defence of the concept [Deutsch, 1998; Lewis, 1986]. Lewis defends the thesis of “modal realism”: that the world we are part of is but one of a plurality of worlds.

³⁶As opposed to consideration of an explicitly hypothetical such ensemble, which can indeed be useful, see Sec. 7.2.

9.2.7 Observations and disproof

Despite the gloomy prognosis given above, there are some specific cases where the existence of a chaotic inflation (multi-domain) type scenario (Sec. 2.6) can be *disproved*. These are firstly when we live in a ‘small universe’ where we have already seen right round the universe (Sec. 4.3.1), for then the universe closes up on itself in a single FL-like domain, so that no further such domains can exist that are causally connected to us in a single connected spacetime. This ‘small universe’ situation is observationally testable (Sec. 4.3.1); its confirmation would disprove the usual chaotic inflationary scenario, but not a truly ‘disconnected’ multiverse proposal, for that cannot be shown to be false by any observation. Neither can it be shown to be true. Secondly, many versions of chaotic inflation, for example those involving Coleman-de Luccia tunneling [Coleman and de Luccia, 1980] from a de Sitter spacetime, demand $k = -1 \Leftrightarrow \Omega_0 < 1$ [Freivogel *et al.*, 2005b; Susskind, 2005]. This requirement is currently marginally disproved by the $2 - \sigma$ bounds on Ω_0 when WMAP observations are combined with the best other available data (Sec. 2.3.7). The best current data is marginally consistent with $k = -1$, but the value indicated most strongly by that data is $k = +1$, indicating finite closed space sections rather than an infinite multiverse such as that advocated by Susskind *et al.* [Freivogel *et al.*, 2005b; Susskind, 2005].

9.2.8 Physical or biological paradigms — Adaptive Evolution?

Given that the multiverse idea must in the end be justified philosophically rather than by scientific testing, is there a philosophically preferable version of the idea? One can suggest there is: greater explanatory power is potentially available by introducing the major constructive principle of biology into cosmology, namely adaptive evolution, which is the most powerful process known that can produce ordered structure where none pre-existed. This is realized in principle in Lee Smolin’s idea (Sec. 9.1.6) of Darwinian adaptation when collapse to black holes is followed by re-expansion, but with an alteration of the constants of physics each time, so as to allow for evolutionary selection towards those regions that produce the maximum number of black holes. The idea needs development, but is very intriguing:

Thesis H4: The underlying physics paradigm of cosmology could be extended to include biological insights. *The dominant paradigm in cosmology is that of theoretical physics. It may be that it will attain deeper explanatory power by embracing biological insights, and specifically that of Darwinian evolution. The Smolin proposal for evolution of populations of expanding universe domains [Smolin, 1992] is an example of this kind of thinking.*

The result is different in important ways from standard cosmological theory precisely because it embodies in one theory three of the major ideas of last century, namely (i) Darwinian evolution of populations through competitive selection, (ii) the evolution of the universe in the sense of major changes in its structure asso-

ciated with its expansion, and (iii) quantum theory, underlying the only partly explicated mechanism supposed to cause re-expansion out of collapse into a black hole. There is a great contrast with the theoretical physics paradigm of dynamics governed simply by variational principles shaped by symmetry considerations. It seems worth pursuing as a very different route to the understanding of the creation of structure.³⁷

9.3 Issue I: Natures of Existence

Underlying all this is the issue of natures of existence, which has a number of aspects, relating from the purely physical to more metaphysical issues.

9.3.1 Physical existence: kinds of matter

Unsolved key issues for physical cosmology relate to what kind of matter and/or fields exist. While we understand matter in the solar system quite well, at present we do not understand most of what exists in the universe at large:

Thesis I1: We do not understand the dominant dynamical matter components of the universe at early or late times. *A key goal for physical cosmology is determining the nature of the inflaton, of dark matter, and of dark energy. Until this is done, the causal understanding of cosmology is incomplete, and in particular the far future fate of the universe is unknown.*

This is the core activity of much work in cosmology at present. Until they are all explicated, cosmology is not properly linked to physics, and the nature of the matter that dominates the dynamics of the universe is unknown. Its explication is surely one of the key concerns of cosmology [Durrer, 2002]. A key requirement is that even if we cannot experimentally verify the proposed nature of the matter, at least it should be physically plausible. This appears not to be the case for some current proposals, e.g. so-called ‘phantom matter’ which has negative kinetic energy terms.

The far future fate of the universe depends crucially on the effective equation of state for dark matter (‘quintessence’). But the problem is that even if we can determine these properties at the present time (for one particular range of parameter values), this does not necessarily guarantee what they will be in the far future (for a quite different range of parameter values that are probably outside the range of possible experimental test). Furthermore adjusting a ‘dark energy’ model to fit the supernova data does not determine the underlying physics. One can fit any monotonic evolution $S(t)$ with a suitable choice of the equation of state function $p(\mu)$. Specifically, for any $S(t)$ and any k we define $\mu(t)$ and $p(t)$ by

$$(41) \quad \kappa\mu(t) = 3 \left[\frac{\dot{S}^2(t)}{S^2(t)} + \frac{k}{S^2(t)} \right], \quad \kappa p(t) = \left[\frac{\dot{S}^2(t)}{S^2(t)} + \frac{k}{S^2(t)} \right] - 2 \frac{\ddot{S}(t)}{S(t)},$$

³⁷Cf. Chapter 13 of Susskind [Susskind, 2005].

then (9), (7) will be exactly satisfied, and we have ‘solved’ the field equations for this arbitrarily chosen monotonic evolution $S(t)$. If we can observationally determine the form of $S(t)$, for example from (m, z) -curves associated with supernovae data, this is essentially how we can then determine that some kind of ‘dark energy’ or ‘quintessence’ is required to give that evolution, and we can find the equation of state implied by eliminating t between these two equations. This is, however, not a *physical* explanation until we have either in some independent experimental test demonstrated that matter of this form exists, or have theoretically shown why this matter or field has the form it does in some more fundamental terms than simply a phenomenological fit. If we assume the matter is a scalar field, the kinetic energy term $\dot{\phi}^2$ implied by (32), (41) may be negative — which is the case for so-called ‘shadow matter’ models proposed recently by some worker. If normal physics criteria are applied, this is a proof that this kind of matter is unphysical, rather than an identification of the nature of the dark energy.

9.3.2 Existence of Infinities

The nature of existence is significantly different if there is a finite amount of matter or objects in the universe, as opposed to there being an infinite quantity in existence. Some proposals claim there may be an infinite number of universes in a multiverse and many cosmological models have spatial sections that are infinite, implying an infinite number of particles, stars, and galaxies. However, infinity is quite different from a very large number! Following David Hilbert [Hilbert, 1964], one can suggest these unverifiable proposals cannot be true: the word ‘infinity’ denotes a quantity or number that can never be attained, and so will never occur in physical reality.³⁸ He states

“Our principal result is that the infinite is nowhere to be found in reality. It neither exists in nature nor provides a legitimate basis for rational thought . . . The role that remains for the infinite to play is solely that of an idea . . . which transcends all experience and which completes the concrete as a totality . . .” [Hilbert, 1964, p. 151].

This suggests “infinity” cannot be arrived at, or realized, in a concrete physical setting; on the contrary, the concept itself implies its inability to be realized!³⁹

Thesis I2: The often claimed physical existence of infinities is questionable. *The claimed existence of physically realized infinities in cosmology or multiverses raises problematic issues. One can suggest they are unphysical; in any case such claims are certainly unverifiable.*

This applies in principle to both small and large scales in any single universe:

³⁸An intriguing further issue is the dual question: Does the quantity zero occur in physical reality? This is related to the idea of physical existence of nothingness, as contrasted with a vacuum [Seife, 2000]. A vacuum is not nothing! (cf. [Susskind, 2005]).

³⁹For a contrasting view, see Bernadete [Bernadete, 1964].

- The existence of a physically existing spacetime continuum represented by a real (number) manifold at the micro-level contrasts with quantum gravity claims of a discrete spacetime structure at the Planck scale, which one might suppose was a generic aspect of fully non-linear quantum gravity theories [Rovelli, 2004]. In terms of physical reality, this promises to get rid of the uncountable infinities the real line continuum engenders in all physical variables and fields.⁴⁰ There is no experiment that can *prove* there is a physical continuum in time or space; all we can do is test space-time structure on smaller and smaller scales, but we cannot approach the Planck scale.
- Infinitely large space-sections at the macro-level raise problems as indicated by Hilbert, and leads to the infinite duplication of life and all events [Ellis and Brundrit, 1979]. We may assume space extends forever in Euclidean geometry and in many cosmological models, but we can never prove that any realised 3-space in the real universe continues in this way — it is an untestable concept, and the real spatial geometry of the universe is almost certainly not Euclidean. Thus Euclidean space is an abstraction that is probably not physically real. The infinities supposed in chaotic inflationary models derive from the presumption of pre-existing infinite Euclidean space sections, and there is no reason why those should necessarily exist. In the physical universe spatial infinities can be avoided by compact spatial sections, resulting either from positive spatial curvature, or from a choice of compact topologies in universes that have zero or negative spatial curvature. Machian considerations to do with the boundary conditions for physics suggest this is highly preferable [Wheeler, 1968]; and if one invokes string theory as a fundamental basis for physics, then ‘dimensional democracy’ suggests the three large spatial dimensions should also be compact, since the small (‘compactified’) dimensions are all taken to be so. The best current data from CBR and other observations (Sec. 2.3.7) indeed suggest $k = +1$, implying closed space sections for the best-fit FL model.
- The existence of an eternal universe implies that an infinite time actually exists, which has its own problems: if an event happens at any time t_0 , one needs an explanation as to why it did not occur before that time (as there was an infinite previous time available for it to occur); and Poincaré eternal return (mentioned in Sec. 6.1) will be possible if the universe is truly cyclic. In any case it is not possible to *prove* that the universe as a whole, or even the part of the universe in which we live, is past infinite; observations cannot do so, and the physics required to guarantee this would happen (if initial conditions were right) is untestable. Even attempting to prove it is future infinite is problematic (we cannot for example guarantee the properties of the

⁴⁰To avoid infinities entirely would require that nothing whatever is a continuum in physical reality (since any continuum interval contains an infinite number of points). Doing without that, conceptually, would mean a complete rewrite of many things. Considering how to do so in a way compatible with observation is in my view a worthwhile project.

vacuum into the infinite future — it might decay into a state corresponding to a negative effective cosmological constant).

- It applies to the possible nature of a multiverse. Specifying the geometry of a generic universe requires an infinite amount of information because the quantities necessary to do so are fields on spacetime, in general requiring specification at each point (or equivalently, an infinite number of Fourier coefficients): they will almost always not be algorithmically compressible. All possible values of all these components in all possible combinations will have to occur in a multiverse in which “all that can happen, does happen”. There are also an infinite number of topological possibilities. This greatly aggravates all the problems regarding infinity and the ensemble. Only in highly symmetric cases, like the FL solutions, does this data reduce to a finite number of parameters, each of which would have to occur in all possible values (which themselves are usually taken to span an infinite set, namely the entire real line). Many universes in the ensemble may themselves have infinite spatial extent and contain an infinite amount of matter, with all the problems that entails. To conceive of physical creation of an infinite set of universes (most requiring an infinite amount of information for their prescription, and many of which will themselves be spatially infinite) is at least an order of magnitude more difficult than specifying an existent infinitude of finitely specifiable objects.

One should note here particularly that problems arise in the multiverse context from the continuum of values assigned by classical theories to physical quantities. Suppose for example that we identify corresponding times in the models in an ensemble and then assume that *all* values of the density parameter and the cosmological constant occur at each spatial point at that time. Because these values lie in the real number continuum, this is a doubly uncountably infinite set of models. Assuming genuine physical existence of such an uncountable infinitude of universes is the antithesis of Occam’s razor. But on the other hand, if the set of realised models is either finite or countably infinite, then almost all possible models are not realised. And in any case this assumption is absurdly unprovable. We can’t observationally demonstrate a single other universe exists [Gardner, 2003], let alone an infinitude. The concept of infinity is used with gay abandon in some multiverse discussions [Knobe *et al.*, 2005], without any concern either for the philosophical problems associated with this statement [Hilbert, 1964], or for its completely unverifiable character. It is an extravagant claim that should be treated with extreme caution.

9.3.3 *The Nature of the Laws of Physics*

Underlying all the above discussion is the basic concept of ordered behaviour of matter, characterized by laws of physics of a mathematical nature that are the

same everywhere in the universe.⁴¹ Three interlinked issues arise.

(i) *What is the ontological nature of the laws of physics: descriptive, just characterizing the way things are, or prescriptive, enforcing them to be this way?* [Carroll, 2004]. If they are descriptive, the issue arising is, *Why does all matter have the same properties wherever it exists in the universe?* Why are all electrons everywhere in the universe identical, if the laws are only descriptive? If they are prescriptive, then matter will necessarily be the same everywhere (assuming the laws themselves are invariable); the issue arising then is, *In what way do laws of physics exist that enforce themselves on the matter in the universe?* Do they for example have an existence in some kind of Platonic space that controls the nature of matter and existence? One can avoid talking about the laws of physics *per se* by instead considering the *space of possibilities* underlying what exists physically, rigorously constraining the possible natures of what actually comes into existence [Ellis, 2004]. This space is more or less uniquely related to the underlying laws in the same way that the space of solutions of differential equations is related to the nature of the equations. This enables one to avoid the issue of the ontology of the laws of physics, but does not solve it.

(ii) *Why are the laws of physics so well explained by mathematical descriptions?* If they are prescriptive, this deep issue might be related to the suggested Platonic nature of the space of mathematical reality [Penrose, 2004]. If they are descriptive, then the mathematical expressions we use to encapsulate them are just a convenient description but do not reflect their ultimate nature. Many writings in physics and cosmology seem to assume that their ultimate existential nature is indeed mathematical — perhaps a confusion of appearance and reality (see Sec. 8.3).

(iii) *Do they pre-exist the universe and control its coming into being, or do they come into being with the universe?* This is where this issue relates deeply to the nature of cosmology, and is clearly related to the other two questions raised above. Many theories of creation of the universe assume that all these laws, or at least a basic subset, pre-exist the coming into being of the physical universe, because they are presumed to underlie the creation process, for example the entire apparatus of quantum field theory is often taken for granted as pre-existing our universe (Sec. 6). This is of course an unprovable proposition

Thesis I3: A deep issue underlying the nature of cosmology is the nature of the laws of physics. *The nature of the possibility space for physical existence is characterized by the laws of physics. However it is unclear if these laws are prescriptive or descriptive; whether they come into being with space-time and matter, or pre-exist them.*

⁴¹The effective laws may vary from place to place because for example the vacuum state varies [Susskind, 2005]; but the fundamental laws that underlie this behaviour are themselves taken to be invariant.

9.3.4 ‘Ultimate Reality’

Philosophers have debated for millennia whether the ultimate nature of existence is purely material, or embodies some form of rationality (‘Logos’) and/or purpose (‘Telos’). *What in the end underlies it all?* Is the ultimate nature of the universe purely material, or does it in some way have an element of the mental? (cf. Sec. 9.1.6). That profound debate is informed by physical cosmology, but cannot be resolved by the physical sciences alone (Sec. 9.1.7). Here, I will make just two comments on this deep issue.

Firstly, even in order to understand just the material world, it can be claimed that one needs to consider forms of existence other than the material only — for example a Platonic world of mathematics and a mental world, both of which can be claimed to exist and be causally effective in terms of affecting the material world [Ellis, 2004; Penrose, 2004]. Our understanding of local causation will be incomplete unless we take them into account.

Secondly, in examining these issues one needs to take into account data about the natures of our existence that come from our daily lives and the broad historical experience of humanity (our experiences of ethics and aesthetics, for example), as well as those discoveries attained by the scientific method. Many writings claim there is no purpose in the universe: it is all just a conglomerate of particles proceeding at a fundamental level in a purposeless and meaningless algorithmic way. But I would reply, the very fact that those writers engage in such discourse undermines their own contention; they ignore the evidence provided by their own actions. There is certainly meaning in the universe to this degree: *the fact they take the trouble to write such contentions is proof that they consider it meaningful to argue about such issues*; and this quality of existence has emerged out of the nature of the physical universe (Sec. 7.3). Indeed the human mind is causally effective in the real physical world precisely through many activities motivated by meanings perceived by the human mind. Any attempt to relate physics and cosmology to ultimate issues must take such real world experience seriously [Ellis, 2005], otherwise it will simply be ignoring a large body of undeniable data. This data does not resolve the ultimate issues, but does indicate dimensions of existence that indeed do occur.

10 CONCLUSION

The physical scale of the Universe is enormous, and the images of distant objects from which we obtain our information are extremely faint. It is remarkable that we are able to understand the Universe as well as we do. An intriguing feature is the way in which the philosophy of cosmology is to a considerable degree shaped by contingent aspects of the nature of the universe — its vast scale (Sec. 4), leading to the existence of visual horizons (Sec. 4.3), and the occurrence of extreme energies in the early universe (Sec. 5), leading to the existence of physical horizons. Philosophical issues arising in relation to cosmology (Sec. 8) would be quite

different if its physical structure were very different. Furthermore in order that philosophical analysis can engage with cosmology in depth, the detailed nature of the relation between observations and theory in cosmology (Sec. 2) is relevant.

10.1 *Are there laws of cosmology?*

As we have discussed in detail, the uniqueness of the universe implies the unique nature of cosmology. We now return to the initial issue, *Are there Laws of the Universe?* (Sec. 3). At one level, the laws of the cosmos are simply the local laws we know and love (e.g. Maxwell's laws, Einstein's field equations) applied to the whole shebang. Of course, there is the problem of extrapolation from the local to the global. But although the extrapolation is bigger in cosmology, it seems not to be different in kind from what we always do in science. In that sense, there are no special laws for the evolution of the universe. But that does not determine the outcome: cosmology needs some prescription of boundary or initial conditions as well, in order to determine the future. Is there a true "Cosmological principle", a law of initial conditions for the universe, that determines this outcome?

The idea of "Laws of initial conditions of the universe" seems not to be a testable idea (Sec. 3). Scientifically, one can only describe what occurred rather than relate it to generic principles, for such principles cannot be tested. In fact any description of boundary or initial conditions for the universe seems to be just that: a description of these conditions, rather than a testable prescription of how they must be. The 'Cosmological Principle' — the universe is necessarily spatially homogeneous and isotropic (Sec. 4.2.2) is of this kind: a description of the way the initial data turned out, rather than a fundamental reason for why this should be so. Justification of this view was based by some workers on a *Copernican Principle* (the assumption we do not live in a privileged place in the universe), strengthened to become a *Cosmological Principle* [Bondi, 1960; Weinberg, 1972; Harrison, 2000]; but this is a philosophical assumption — essentially, a claim that the universe must have very special initial conditions — which may or may not be true, and does not attempt a physical explanation. This kind of argument is out of fashion at present, because we now prefer generality to speciality and physical argumentation to geometrical prescription; but it was previously strongly proposed (e.g. [Weinberg, 1972], pp. 407-412). The tenor of philosophical argument has changed.

Nevertheless there is one kind of Law of the Universe one might propose, following McCrea [McCrea, 1970]: namely an "Uncertainty principle in cosmology", dual to the uncertainty principle in quantum theory. Uncertainty applies on the largest scale, as we have discussed above in some detail, and also on the smallest, where it is a profound feature of quantum theory. Its basis is very different in the two cases, on the one hand (in quantum theory) being ontological in nature, on the other (in cosmology) being epistemological in nature.⁴² Nevertheless it is

⁴²Assuming that quantum uncertainty is indeed ontological rather than epistemological. One should however keep an open mind on this: just because it is the current dogma does not

a key aspect of our relation to the cosmos, so that (following McCrea) we might perhaps formalize it in order to emphasize its centrality to the relation between cosmology and philosophy:

Thesis of Uncertainty: Ultimate uncertainty is a key aspect of cosmology. *Scientific exploration can tell us much about the universe but not about its ultimate nature, or even much about some of its major geometrical and physical characteristics. Some of this uncertainty may be resolved, but much will remain. Cosmological theory should acknowledge this uncertainty.*

10.2 *What can we truly claim*

Cosmology considers questions of physical origins in the uniquely existing physical universe (Sec. 6) which provides the context of our existence (Sec. 7, Sec. 9.1). These questions can be extended to include ultimate issues if we so desire (Sec. 8.2), but physical theory cannot resolve them (Sec. 9.1.7). In the end, there are a variety of mysteries underlying the existence and nature of the universe (Sec. 9.3). The scientific study of cosmology can help illuminate their nature, but cannot resolve them.

As well as celebrating the achievements of cosmology, one should fully take into account the limits and problems considered in this chapter, and not claim for scientific cosmology more than it can actually achieve or more certainty than is in fact attainable. Such claims will in the long term undermine cosmology's legitimate status as a project with solid scientific achievements to its name. That status can be vigorously defended as regards the 'Standard Model' of cosmology (Sec. 2.8), provided this standard model is characterized in conservative terms so that it is not threatened by relatively detailed shifts in theory or data that do not in fact threaten the core business of cosmology. Further, this defence must take adequate cognisance of the difficult philosophical issues that arise if one pushes the explanatory role of cosmological theory to its limits (Sec. 6); for example one should not make too strong *scientific* claims in regard to the possible existence of multiverses (Sec. 9.2); philosophically based plausibility arguments for them are fine, if identified as such. Cosmology is not well served by claims that it can achieve more explanatory power than is in fact attainable, or by statements that its claims are verified when in fact the requisite evidence is unavailable, and in some cases must forever remain so.

ACKNOWLEDGEMENT

I thank Bill Stoeger, Martin Bojowald, Malcolm MacCallum, Henk van Elst, Jeremy Butterfield, and John Earman for useful comments that have improved this article.

necessarily mean it is true.

Abbreviations used:

- CBR: Cosmic Blackbody Radiation, p. 1192
 CDM: Cold Dark Matter, p. 1201
 EFE: Einstein Field Equations, p. 1185
 FL: Friedmann-Lemaître (universe models), p. 1187
 HBB: Hot Big Bang, p. 1191
 LSS: Last Scattering surface, p. 1192
 RW: Robertson-Walker (geometry), p. 1186
 SAP: Strong Anthropic Principle, p. 1254
 WAP: Weak Anthropic Principle, p. 1253

BIBLIOGRAPHY

- [Allday, 2002] J. Allday. *Quarks, Leptons and the Big Bang*. Institute of Physics, 2002.
- [Balashov, 1991] Y. V. Balashov. Resource letter Ap-1: The anthropic principle. *American Journal of Physics*, 54:1069–1076, 1991.
- [Barbour and Pfister, 1995] J. Barbour and H. Pfister. *Mach's Principle: From Newton's Bucket to Quantum Gravity*. Birkhäuser, Basel, 1995.
- [Bardeen, 1980] J. M. Bardeen. Gauge-invariant cosmological perturbations. *Phys. Rev.*, D22:1882-1905, 1980.
- [Barnett et al., 1996] R. M. Barnett et al. (Particle Data Group) Particle physics summary. a digest of the 1996 Review of Particle Physics. *Reviews of Modern Physics*, 68:611–732, 1996.
- [Barrow and Tipler, 1984] J. Barrow and F. Tipler. *The Cosmological Anthropic Principle*. Oxford University Press, 1984.
- [Barrow, 2003] J. D. Barrow. *The Constants of Nature*. Vintage, 2003.
- [Bennet et al., 2003] C. L. Bennet et al. (the WMAP team). First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Preliminary maps and basic results. *Astrophysical Journal Supplement Series*, 148:1–27, 2003. astro-ph/0302207.
- [Berendzen et al., 1976] R. Berendzen, R. Hart and D. Seeley. *Man Discovers the Galaxies*. Science History Publications, New York, 1976.
- [Bernadete, 1964] J. A. Bernadete. *Infinity, An Essay in Metaphysics*. Clarendon Press, Oxford, 1964.
- [Blau and Guth, 1987] S. K. Blau and A. H. Guth. Inflationary cosmology. In Hawking, S. W. and Israel, W., editors, *300 Years of Gravitation*, page 524. Cambridge University Press, 1987.
- [Bojowald, 2001] M. Bojowald. Absence of singularity in loop quantum cosmology. *Physical Review Letters*, 86:5227–5239, 2001.
- [Bondi, 1947] H. Bondi. Spherically symmetric models in general relativity. *Monthly Notice of the Royal Astronomical Society*, 107:410, 1947.
- [Bondi, 1960] H. Bondi. *Cosmology*. Cambridge University Press, Cambridge, 1960.
- [Bonnor and Ellis, 1986] W. B. Bonnor and G. F. R. Ellis. Observational homogeneity of the universe. *Monthly Notice of the Royal Astronomical Society*, 218:605–614, 1986.
- [Bostrom, 2002] N. Bostrom. *Anthropic bias: Observation selection effects in science and philosophy*. Routledge, New York, 2002.
- [Bothun, 1998] G. Bothun. *Modern Cosmological Observations and Problems*. Taylor and Francis, 1998.
- [Brading and Castellani, 2006] K. Brading and E. Castellani. This vol, Ch. 13, 2006.
- [Brans and Dicke, 1961] C. Brans and R. H. Dicke. Mach's principle and a relativistic theory of gravitation. *Physical Review*, 124:925, 1961.
- [Carr and Rees, 1979] B. J. Carr and M. J. Rees. The anthropic principle and the structure of the physical world. *Nature*, 278:605–612, 1979.
- [Carroll, 2004] J. Carroll, editor. *Readings on Laws of Nature*. University of Pittsburgh Press, 2004.
- [Carroll and Chen, 2005] S. M. Carroll and J. Chen. Does inflation provide natural initial conditions for the universe? gr-qc/0505037, 2005.

- [Challinor and Lasenby, 1998] A. D. Challinor and A. N. Lasenby. Cosmic microwave background anisotropies in the CDM model: A covariant and gauge-invariant approach. *Physical Review D*, 58. astro-ph/9804301, 1998.
- [Coleman and de Luccia, 1980] S. Coleman and F. de Luccia. Gravitational effects on and of vacuum decay. *Physical Review D*, 21:3305–3315, 1980.
- [Copeland *et al.*, 2005] E. J. Copeland, J. E. Lidsey, and S. Mizuno. Correspondence between loop-inspired and braneworld cosmology. gr-qc/0510022, 2005.
- [Cowie and Songaila, 1995] L. Cowie and A. Songaila. Astrophysical limits on the evolution of dimensionless physical constants over cosmological time. *Astrophysics Journal*, 453:596–598, 1995.
- [Davies, 1974] P. C. W. Davies. *The physics of time asymmetry*. Surrey University Press, London, 1974.
- [Davies, 1982] P. C. W. Davies. *The Accidental Universe*. Cambridge University Press, 1982.
- [Davies, 1987] P. C. W. Davies. *The cosmic Blueprint*. Heinemann, London, 1987.
- [Davis and Lineweaver, 2004] T. M. Davis. and C. H. Lineweaver. Expanding confusion: common misconceptions of cosmological horizons and the superluminal expansion of the universe. *Publications of the Astronomical Society of Australia*, 21:97–109, 2004. astro-ph/0310808.
- [de Vaucouleurs, 1970] G. de Vaucouleurs. The case for a hierarchical cosmology. *Science*, 167:1203–13, 1970.
- [Deutsch, 1998] D. Deutsch. *The Fabric of reality: The science of parallel universes - and its implications*. Penguin, 1998.
- [Dicke and Peebles, 1979] R. H. Dicke. and P. J. E. Peebles. The big bang cosmology—enigmas and nostrums. In Hawking, S. W. and Israel, W., editors, *General Relativity: An Einstein Centenary Survey*, pages 504–517. Cambridge University Press, Cambridge, 1979.
- [Dickson, 2006] M. I. Dickson. this vol , Ch.4, 2006.
- [Dirac, 1938] P. A. M. Dirac. New basis for cosmology. *Proceedings of the Royal Society of London*, 165:199–208, 1938.
- [Disney, 1976] M. J. Disney. Visibility of galaxies. *Nature*, 263:573–575, 1976.
- [Dodelson, 2003] S. Dodelson. *Modern Cosmology*. Academic Press, 2003.
- [Durrer, 2002] R. Durrer. Frontiers of the universe: What do we know and what do we understand? astro-ph/0205101, 2002.
- [Earman, 1987] J. Earman. The SAP also rises: A critical examination of the anthropic principle. *American Philosophical Quarterly*, 24:307–325, 1987.
- [Earman, 1999] J. Earman. The Penrose-Hawking singularity theorems: History and implications. In Goenner, H., Renn, J., and Sauer, T., editors, *The Expanding Worlds of General Relativity*. Birkhäuser, 1999.
- [Earman, 2001] J. Earman. Lambda: The constant that would not die. *Archive for History of Exact Sciences*, 55:189–220, 2001.
- [Earman, 2003] J. Earman. The cosmological constant, the fate of the universe, unimodular gravity, and all that. *Studies in History and Philosophy of Modern Physics*, 34:559–577, 2003.
- [Earman and Mosterin, 1999] J. Earman and J. Mosterin. A critical look at inflationary cosmology. *Philosophy of Science*, 66:1–49, 1999.
- [Ehlers, 1993] J. Ehlers. Contributions to the relativistic mechanics of continuous media. *General Relativity and Gravitation*, 25:1225–1266, 1993.
- [Ehlers *et al.*, 1968] J. Ehlers, P. Geren, and R. K. Sachs. Isotropic solutions of the Einstein-Liouville equations. *Journal of Mathematical Physics*, 9:1344–1349, 1968.
- [Ehlers and Rindler, 1989] J. Ehlers and W. Rindler. A phase-space representation of Friedmann-Lemaître universes containing both dust and radiation and the inevitability of a big bang. *Monthly Notice of the Royal Astronomical Society*, 238:503–521, 1989.
- [Eisenstein *et al.*, 2005a] D. J. Eisenstein *et al.* Detection of the baryon acoustic peak in the large scale correlation function of galaxies, 2005. astro-ph/0501171.
- [Ellis, 1967] G. F. R. Ellis. The dynamics of pressure-free matter in general relativity. *Journal of Mathematical Physics*, 8:1171–1194, 1967.
- [Ellis, 1971a] G. F. R. Ellis. Relativistic cosmology. In Sachs, R. K., editor, *General Relativity and Cosmology, Proceedings of the XLVII Enrico Fermi Summer School*. Academic Press, New York, 1971.
- [Ellis, 1971b] G. F. R. Ellis. Topology and cosmology. *General Relativity and Gravitation*, 2:7–21, 1971.

- [Ellis, 1975] G. F. R. Ellis. Cosmology and verifiability. *Quarterly Journal of the Royal Astronomical Society*, 16:245–264, 1975.
- [Ellis, 1980] G. F. R. Ellis. Limits to verification in cosmology. *Ann New York Acad Sci*, 336:130–160, 1980.
- [Ellis, 1984] G. F. R. Ellis. Relativistic cosmology: its nature, aims and problems. In *et al.*, B. B., editor, *General Relativity and Gravitation*, pages 215–288. Reidel, 1984.
- [Ellis, 1989] G. F. R. Ellis. A history of cosmology 1917–1955. In Howard, D. and Stachel, J., editors, *Einstein and the History of General Relativity*, volume 1 of *Einstein Study Series*, pages 367–431. Birkhäuser, Boston, 1989.
- [Ellis, 1990] G. F. R. Ellis. Innovation resistance and change: the transition to the expanding universe. In Bertotti, B., Balbinto, R., Bergia, S., and Messina, A., editors, *Modern Cosmology in Retrospect*, pages 97–114. Cambridge University Press, 1990.
- [Ellis, 1991] G. F. R. Ellis. Major themes in the relation between philosophy and cosmology. *Mem Ital Ast Soc*, 62:553–605, 1991.
- [Ellis, 2002] G. F. R. Ellis. Cosmology and local physics. *New Astronomy Reviews*, 46:645–658, 2002. gr-qc/0102017.
- [Ellis, 2004] G. F. R. Ellis. True complexity and its associated ontology. In Barrow, J. D., Davies, P. C., , and Jr., C. L. H., editors, *Science and Ultimate Reality: Quantum Theory, Cosmology and Complexity, honoring John Wheeler's 90th birthday*. Cambridge University Press, 2004.
- [Ellis, 2005] G. F. R. Ellis. Physics, complexity, and causality. *Nature*, 435:743, 2005.
- [Ellis and Baldwin, 1984] G. F. R. Ellis. and J. Baldwin. On the expected anisotropy of radio source counts. *Monthly Notice of the Royal Astronomical Society*, 206:377–381, 1984.
- [Ellis and Brundrit, 1979] G. F. R. Ellis and G. B. Brundrit. Life in the infinite universe. *Quarterly Journal of the Royal Astronomical Society*, 20:37–41, 1979.
- [Ellis and Bruni, 1989] G. F. R. Ellis and M. Bruni. A covariant and gauge-free approach to density fluctuations in cosmology. *Physical Review D*, 40:1804–1818, 1989.
- [Ellis and Buchert, 2006] G. F. R. Ellis and T. R. Buchert. The universe seen at different scales. To appear, *Physics Letter A*, special Einstein issue, 2006. gr-qc/0506106.
- [Ellis and King, 1974] G. F. R. Ellis and A. R. King. Was the big bang a whimper? *Communications in Mathematical Physics*, 38:119–165, 1974.
- [Ellis *et al.*, 2004] G. F. R. Ellis, U. Kirchner, and W. Stoeger. Multiverses and physical cosmology. *Monthly Notice of the Royal Astronomical Society*, 347:921–936, 2004.
- [Ellis and Maartens, 2004] G. F. R. Ellis and R. Maartens. The emergent universe: inflationary cosmology with no singularity and no quantum gravity era. *Classical and Quantum Gravity*, 21:223–232, 2004.
- [Ellis *et al.*, 1978] G. F. R. Ellis, R. Maartens, and S. D. Nel. The expansion of the universe. *Monthly Notice of the Royal Astronomical Society*, 184:439–465, 1978.
- [Ellis and Madsen, 1991] G. F. R. Ellis and M. Madsen. Exact scalar field cosmologies. *Classical and Quantum Gravity*, 8:667–676, 1991.
- [Ellis *et al.*, 2002a] G. F. R. Ellis, P. McEwan, W. Stoeger, and P. Dunsby. Causality in inflationary universes with positive spatial curvature. *General Relativity and Gravitation*, 34:1461–1481, 2002. gr-qc/0109024.
- [Ellis *et al.*, 1985] G. F. R. Ellis, S. D. Nel, W. Stoeger, R. Maartens, and A. P. Whitman. Ideal observational cosmology. *Physics Reports*, 124:315–417, 1985.
- [Ellis *et al.*, 1984] G. F. R. Ellis, J. J. Perry, and A. Sievers. Cosmological observations of galaxies: the observational map. *Astronomical Journal*, 89:1124–1154, 1984.
- [Ellis and Schreiber, 1986] G. F. R. Ellis and G. Schreiber. Observational and dynamic properties of small universes. *Physics Letters A*, 115:97–107, 1986.
- [Ellis and Sciamia, 1972] G. F. R. Ellis and D. W. Sciamia. Global and non-global problems in cosmology. In O’Raifeartaigh, L., editor, *General Relativity*, pages 35–59. Oxford University Press, 1972.
- [Ellis *et al.*, 2002b] G. F. R. Ellis, W. Stoeger, P. McEwan, and P. Dunsby. Dynamics of inflationary universes with positive spatial curvature. *General Relativity and Gravitation*, 34:1445–1459, 2002. gr-qc/0109023.
- [Ellis and Stoeger, 1987] G. F. R. Ellis and W. R. Stoeger. The Fitting problem in cosmology. *Classical and Quantum Gravity*, 4:1679–1690, 1987.
- [Ellis and Stoeger, 1988] G. F. R. Ellis and W. R. Stoeger. Horizons in inflationary universes. *Classical and Quantum Gravity*, 5:207, 1988.

- [Ellis and Uzan, 2005] G. F. R. Ellis and J.-P. Uzan. ‘c’ is the speed of light, isn’t it? *American Journal of Physics*, 73:240–247, 2005.
- [Ellis and van Elst, 1999a] G. F. R. Ellis and H. van Elst. Cosmological models (cargese lectures 1998). In Lachièze-Ray, M., editor, *Theoretical and Observational Cosmology*, volume 541, pages 1–116. Kluwer, Nato Series C: Mathematical and Physical Sciences, 1999. gr-qc/9812046.
- [Ellis and van Elst, 1999b] G. F. R. Ellis and H. van Elst. Deviation of geodesics in FLRW spacetime geometries. In Harvey, A., editor, *On Einstein’s Path: Essays in Honour of Engelbert Schucking*, pages 203–226. Springer, 1999.
- [Ellis and Williams, 2000] G. F. R. Ellis and R. M. Williams. *Flat and Curved Spacetimes*. Cambridge University Press, 2000.
- [Fabian, 1989] A. C. Fabian, editor. *Origins*. Cambridge University Press, 1989.
- [Freivogel *et al.*, 2005a] B. Freivogel, M. Kleban, M. R. Martinez, and L. Susskind. Observational consequences of a landscape. hep-th/0505232, 2005.
- [Freivogel *et al.*, 2005b] B. Freivogel, M. Kleban, M. R. Martinez, and L. Susskind. Observational consequences of a landscape. hep-th/0505232, 2005.
- [Gardner, 1986] M. Gardner. WAP, SAP, PAP, and FAP. *The New York Review of Books*, 23(8):22–25, 1986.
- [Gardner, 2003] M. Gardner. *Are Universes Thicker than Blackberries?* W W Norton, 2003.
- [Garrett and Coles, 1993] A. J. M. Garrett and P. Coles. Bayesian inductive inference and the anthropic cosmological principle. *Comments Astrophys*, 17:23–47, 1993.
- [Ghosh *et al.*, 1998] A. Ghosh, G. Pollifrone, and G. Veneziano. Quantum fluctuations in open pre-big bang cosmology. *Physics Letters B*, 440:20–27, 1998.
- [Gibbons *et al.*, 1983] G. W. Gibbons, S. W. Hawking, and S. T. C. Siklos. *The Very Early Universe*. Cambridge University Press, 1983.
- [Gibbons *et al.*, 2003] G. W. Gibbons, E. P. S. Shellard, and S. J. Rankin, editors. *The Future of Theoretical Physics and Cosmology: Celebrating Stephen Hawking’s 60th Birthday*. Cambridge University Press, 2003.
- [Goodman, 1995] J. Goodman. Geocentrism reexamined. *Physical Review D*, 52:1821–1827, 1995.
- [Gott and Li, 1998] J. R. Gott and L.-X. Li. Can the universe create itself? *Physical Review D*, 58:023501, 1998. astro-ph/9712344.
- [Gribbin and Rees, 1991] J. Gribbin and M. Rees. *Cosmic Coincidences*. Black Swan, 1991.
- [Grunbaum, 1989] A. Grunbaum. The pseudo problem of creation. *Philosophy of Science*, 56:373–394, 1989. reprinted in [Leslie, 1998].
- [Grunbaum, 2004] A. Grunbaum. The poverty of theistic cosmology. *British Journal of Philosophy of Science*, 55:561–614, 2004.
- [Gunn *et al.*, 1978] J. E. Gunn, M. S. Longair and M. J. Rees. *Observational cosmology: Saas-Fee Course 8*. Geneva Observatory, Switzerland, 1978.
- [Guth, 1997] A. H. Guth. *The Inflationary Universe: The Quest for a new Theory of Cosmic Origins*. Addison Wesley, 1997.
- [Guth, 1981] A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness. *Physical Review D*, 23:347, 1981.
- [Guth, 2001] A. H. Guth. Eternal inflation. In *Proceedings of “Cosmic Questions” Meeting*, volume 95. The New York Academy of Sciences Press, 2001. astro-ph/0101507.
- [Harris *et al.*, 1998] W. E. Harris, P. R. Durell, M. J. Pierce, and J. Seckers. Constraints in the Hubble constant from observations of the brightest red giant stars in a Virgo-cluster galaxy. *Nature*, 395:45, 1998.
- [Harrison, 2000] E. R. Harrison. *Cosmology: The Science of the Universe*. Cambridge University Press, Cambridge, 2 edition, 2000.
- [Hartle, 2004] J. B. Hartle. Anthropic reasoning and quantum cosmology. In Allen, R. and Pope, C., editors, *The New Cosmology*. AIP, 2004. gr-qc/0406104.
- [Harvey, 2006] J. Harvey. this vol, Ch. 14, 2006.
- [Harwit, 1984] M. Harwit. *Cosmic Discovery*. MIT Press, 1984.
- [Hawking, 1987] S. W. Hawking. Quantum cosmology. In Hawking, S. W. and Israel, W., editors, *300 Years of Gravitation*, page 631. Cambridge University Press, Cambridge, 1987.
- [Hawking, 1993] S. W. Hawking. *On the Big Bang and Black Holes*. World Scientific, 1993.
- [Hawking and Ellis, 1973] S. W. Hawking and G. F. R. Ellis. *The Large Scale Structure of Space-time*. Cambridge University Press, 1973.

- [Hilbert, 1964] D. Hilbert. On the infinite. In Benacerraf, P. and Putnam, H., editors, *Philosophy of Mathematics*, pages 134–151. Englewood Cliff, N J.: Prentice Hall, 1964.
- [Hogan, 1999] C. J. Hogan. Where the baryons are. astro-ph/9912107, 1999.
- [Hogan, 2003] C. J. Hogan. Why the universe is just so. *Reviews of Modern Physics*, 72:1149–1161, 2003.
- [Hollands and Wald, 2002] S. Hollands and R. M. Wald. An alternative to inflation. *Gen Rel. Grav.*, 34:2043–2055, 2002.
- [Hoyle, 1960] F. Hoyle. Cosmological tests of gravitational theories. In *Rendiconti Scuola Enrico Fermi XX Corso*, pages 141–. Academic Press, New York, 1960.
- [Hu and Sugiyama, 1995a] W. Hu and N. Sugiyama. Anisotropies in the CMB: An analytic approach. *Astrophysical Journal*, 444:489–506, 1995.
- [Hu and Sugiyama, 1995b] W. Hu and N. Sugiyama. Toward understanding CMB anisotropies and their implications. *Physical Review D*, 51:2599–2630, 1995.
- [Hubble, 1936] E. Hubble. *The Realm of the Nebulae*. Yale University Press, 1936.
- [Hume, 1993] D. Hume. *Dialogues and Natural History of Religion*. World Classics. Oxford University Press, 1993.
- [Isham, 1997] C. J. Isham. *Lectures on Quantum Theory: Mathematical and Structural Foundations*. Imperial College Press, London, 1997.
- [Jain and Dev, 2005] D. Jain and A. Dev. Age of high redshift objects - a litmus test for dark energy models. astro-ph/0509212, 2005.
- [Kamionkowski and Loeb, 1997] M. Kamionkowski and A. Loeb. Getting around cosmic variance. *Physical Review D*, 56:4511, 1997.
- [Kantowski, 1998] R. Kantowski. The effects of inhomogeneities on evaluating the mass parameter ω_m and the cosmological constant λ . *Astrophysical Journal*, 507:483–496, 1998.
- [Kantowski et al., 1995] R. Kantowski, T. Vaughan, and D. Branch. The effects of inhomogeneities on evaluating the deceleration parameter. *Astrophysical Journal*, 447:35–42, 1995.
- [Khoury et al., 2001] J. Khoury, B. A. Ovrut, P. J. Steinhardt, and N. Turok. The ekpyrotic universe: Colliding branes and the origin of the hot big bang. *Physical Review D*, 64:123522, 2001.
- [Knobe et al., 2005] J. Knobe, K. D. Olum, and A. Vilenkin. Philosophical implications of inflationary cosmology. phys/0320071, 2005.
- [Kolb and Turner, 1990] E. W. Kolb and M. S. Turner. *The Early Universe*. Addison Wesley, 1990.
- [Kragh, 1996] H. Kragh. *Cosmology and Controversy*. Princeton University Press, Princeton, 1996.
- [Kristian and Sachs, 1966] J. Kristian and R. K. Sachs. Observations in cosmology. *Astrophysical Journal*, 143:379, 1966.
- [Kuhn, 1977] T. Kuhn. Objectivity, value judgment, and theory choice. In *The Essential Tension*, pages 320–39. University of Chicago Pressm 1977.
- [Lachièze et al., 1995] M. M. R. Lachièze and J. P. Luminet. Cosmic topology. *Physics Reports*, 254:135, 1995.
- [Lakatos, 1980] I. Lakatos. *The Methodology of Scientific Research Programmes*. Cambridge University Press, 1980.
- [Landsman, 2006] N. P. Landsman. this vol, ch.5, 2006.
- [Langlois and Vernizzi, 2005] D. Langlois and F. Vernizzi. Conserved non-linear quantities in cosmology. astro-ph/0509078, 2005.
- [Lemaître, 1931] G. Lemaître. The beginning of the world from the point of view of quantum theory. *Nature*, 127:706, 1931.
- [Leslie, 1989] J. Leslie. *Universes*. Routledge, 1989.
- [Leslie, 1994] J. Leslie. Cosmology: a philosophical survey. *Philosophia*, 24:3–27, 1994.
- [Leslie, 1998] J. Leslie, editor. *Modern Cosmology and Philosophy*. Prometheus Books, 1998.
- [Lewis, 1986] D. K. Lewis. *On the Plurality of Worlds*. Basil Blackwell, Oxford, 1986.
- [Liddle, 2004] A. R. Liddle. How many cosmological parameters? astro-ph/0401198, 2004.
- [Lidsey et al., 1997] J. E. Lidsey, A. R. Liddle, E. W. Kolb, E. J. Copeland, T. Barriero, and M. Abney. Reconstructing the inflation potential—an overview. *Reviews of Modern Physics*, 69,:373–410, 1997.
- [Linde, 1990] A. D. Linde. *Particle Physics and Inflationary Cosmology*. Harwood Academic, 1990.

- [Lineweaver and Davis, 2005] C. H. Lineweaver and T. M. Davis. Misconceptions about the big bang. *Scientific American*, 2005.
- [Luminet, 2005] J.-P. Luminet. A cosmic hall of mirrors. *Physics World*, 18:23–28, 2005.
- [Luminet *et al.*, 2003] J.-P. Luminet, J. R. Weeks, A. Raizuolo, R. Lehoucq, and J. P. Uzan. Dodecahedral space topology as an explanation for weak wide-angle temperature correlations in the cosmic microwave background. *Nature*, 425:593–595, 2003.
- [Madsen and Ellis, 1988] M. S. Madsen and G. F. R. Ellis. Evolution of Ω in inflationary universes. *Monthly Notice of the Royal Astronomical Society*, 234:67–77, 1988.
- [Malament, 2006] D. Malament. This vol ch. 3, 2006.
- [Matravers *et al.*, 1995] D. R. Matravers, G. F. R. Ellis, and W. R. Stoeger. Complementary approaches to cosmology: Relating theory and observations. *Quarterly Journal of the Royal Astronomical Society*, 36:29–45, 1995.
- [McCrea, 1953] W. H. McCrea. Cosmology. *Reports on Progress in Physics*, 16:321–363, 1953.
- [McCrea, 1960] W. H. McCrea. The interpretation of cosmology. *Nature*, 186:1035, 1960.
- [McCrea, 1970] W. H. McCrea. A philosophy for big bang cosmology. *Nature*, 228:21, 1970.
- [Meyer, 1994] D. M. Meyer. A distant space thermometer. *Nature*, 371:13, 1994.
- [Misner, 1968] C. W. Misner. The isotropy of the universe. *Astrophysical Journal*, 151:431–458, 1968.
- [Misner, 1969] C. W. Misner. Mixmaster universe. *Physical Review Letters*, 22:1071–1074, 1969.
- [Morowitz, 2002] L. Morowitz. *The emergence of everything: How the world came to be complex*. Oxford, 2002.
- [Mouri and Taniguchi, 2005] H. Mouri and Y. Taniguchi. Hierarchical object formation in the peculiar velocity field. astro-ph/0508069, 2005.
- [Mulryne *et al.*, 2005] D. J. Mulryne, R. Tavakol, J. E. Lidsey, and G. F. R. Ellis. An emergent universe from a loop. *Physical Review D*, 71:123512, 2005.
- [Munitz, 1962] M. K. Munitz. The logic of cosmology. *British Journal for the Philosophy of Science*, 13:34–50, 1962.
- [Mustapha *et al.*, 1998] N. Mustapha, C. Hellaby, and G. F. R. Ellis. Large scale inhomogeneity vs source evolution: Can we distinguish them? *Monthly Notice of the Royal Astronomical Society*, 292:817–830, 1998.
- [Nilsson *et al.*, 1991] J. S. Nilsson *et al.*, editors. *The Birth and Early Evolution of our Universe: Proceedings of Nobel Symposium 19*. World Scientific, 1991.
- [North, 1965] J. D. North. *The Measure of the Universe*. Oxford University Press, 1965.
- [Padmanbhan, 1993] T. Padmanbhan. *Structure formation in the universe*. Cambridge University Press, 1993.
- [Partridge, 1995] B. Partridge. *3K: The Cosmic Microwave Background Radiation*. Cambridge University Press, 1995.
- [Peacock, 1999] J. A. Peacock. *Cosmological Physics*. Cambridge University Press, 1999.
- [Peebles, 1971] P. J. E. Peebles. *Physical Cosmology*. Princeton University Press, Princeton, 1971.
- [Peebles, 1993a] P. J. E. Peebles. *The Large-Scale Structure of the Universe*. Princeton University Press, 1993.
- [Peebles, 1993b] P. J. E. Peebles. *Principles of Physical Cosmology*. Princeton University Press, 1993.
- [Peebles *et al.*, 1991] P. J. E. Peebles, D. N. Schramm, E. L. Turner, and R. G. Kron. The case for the relativistic hot big bang cosmology. *Nature*, 352:769–776, 1991.
- [Penrose, 1989a] R. Penrose. Difficulties with inflationary cosmology. In Fergus, E., editor, *14th Texas Symposium Relativistic Astrophysics*. New York Academy of Science, New York, 1989.
- [Penrose, 1989b] R. Penrose. *The Emperor's New Mind*. Oxford University Press, 1989.
- [Penrose, 2004] R. Penrose. *The Road to Reality*. Jonathan Cape, 2004.
- [Perez *et al.*, 2005] A. Perez, H. Sahlmann, and D. Sudarsky. On the quantum origin of the seeds of cosmic structure. gr-qc/0508100, 2005.
- [Perkins, 2005] D. Perkins. *Particle Astrophysics*. Oxford, 2005.
- [Perlmutter *et al.*, 1998] S. Perlmutter *et al.* Discovery of a supernova explosion at half the age of the universe. *Nature*, 391:51, 1998.
- [Pettini, 1999] M. Pettini. Element abundances at high redshifts. astro-ph/9902173, 1999.
- [Rees, 1995] M. J. Rees. *Perspectives in Astrophysical Cosmology*. Cambridge University Press, 1995.

- [Rees, 1999] M. J. Rees. *Just six numbers: The Deep Forces that shape the universe*. Weidenfeld and Nicholson, London, 1999.
- [Rees, 2003] M. J. Rees. *Our Cosmic Habitat*. Princeton University Press, Princeton, 2003.
- [Rindler, 1956] W. Rindler. Visual horizons in world models. *Monthly Notice of the Royal Astronomical Society*, 116:662, 1956.
- [Rindler, 2001] W. Rindler. *Relativity: Special, General, and Cosmological*. Oxford University Press, 2001.
- [Robertson, 1933] H. P. Robertson. Relativistic cosmology. *Reviews of Modern Physics*, 5:62–90, 1933.
- [Robertson, 1935] H. P. Robertson. Kinematics and world structure. *Astrophysical Journal*, 82:284–301, 1935.
- [Rothman and Ellis, 1993] T. Rothman and G. F. R. Ellis. Lost horizons cosmology. *American Journal of Physics*, 61:883–893, 1993.
- [Rovelli, 2004] C. Rovelli. *Quantum Gravity*. Cambridge University Press, 2004.
- [Rovelli, 2006] C. Rovelli. This vol ch.12, 2006.
- [Rugh and Zinkernagel, 2002] S. E. Rugh and H. Zinkernagel. The quantum vacuum and the cosmological constant problem. *Studies in History and Philosophy of Modern Physics*, 33:663–705, 2002.
- [Saini *et al.*, 2000] T. D. Saini, S. Raychaudhury, V. Sahni, and A. A. Starobinsky. Reconstructing the cosmic equation of state from supernova distances. *Physical Review Letters*, 85:1162–1165, 2000.
- [Sandage, 1961] A. Sandage. The ability of the 200-inch telescope to distinguish between selected world models. *Astrophysical Journal*, 133:355–392, 1961.
- [Sandage *et al.*, 1993] A. Sandage, R. G. Kron, and M. S. Longair. *The Deep Universe: Saas-Fee Advanced Course 23*. Springer, Berlin, 1993.
- [Schneider *et al.*, 1992] P. Schneider, J. Ehlers, and E. E. Falco. *Gravitational Lenses*. Springer, Berlin, 1992.
- [Schramm and Turner, 1998] D. N. Schramm and M. S. Turner. Big-bang nucleosynthesis enters the precision era. *Reviews of Modern Physics*, 70:303–318, 1998.
- [Schrödinger, 1956] E. Schrödinger. *Expanding Universes*. Cambridge University Press, Cambridge, 1956.
- [Sciama, 1971] D. W. Sciama. Astrophysical cosmology. In Sachs, R. K., editor, *General Relativity and Cosmology*, pages 183–236. Academic Press, New York, 1971.
- [Scott, 2005] D. Scott. The standard cosmological model. astro-ph/0510731, 2005.
- [Seife, 2000] C. Seife. *Zero: the Biography of a Dangerous Idea*. Penguin, 2000.
- [Seife, 2003] C. Seife. *Alpha and Omega: The Search for the Beginning and end of the Universe*. Bantam Books, London, 2003.
- [Seife, 2004] C. Seife. Physics enters the twilight zone. *Science*, 305:464–466, 2004.
- [Sievers *et al.*, 2005b] J. L. Sievers *et al.* Implications of the cosmic background imager polarization data, 2005. astro-ph/0509203.
- [Sigurdson and Furlanetto, 2005] K. Sigurdson and S. R. Furlanetto. Measuring the primordial deuterium abundance during the cosmic dark ages. astro-ph/0505173.
- [Silk, 1997] J. Silk. *A Short History of the Universe*. Scientific American Library, 1997.
- [Silk, 2001] J. Silk. *The Big Bang*. Freeman, 3 edition, 2001.
- [Silk, 2005] J. Silk. *On the shores of the unknown: A short history of the universe*. Cambridge University Press, 2005.
- [Smith, 1982] R. W. Smith. *The Expanding Universe: Astronomy's Great Debate*. Cambridge University Press, 1982.
- [Smolin, 1992] L. Smolin. Did the universe evolve? *Classical and Quantum Gravity*, 9:173–191, 1992.
- [Spergel *et al.*, 2003] D. N. Spergel *et al.* First year Wilkinson microwave anisotropy probe (WMAP) observations: Determination of cosmological parameters. *Astrophysical Journal Supplement Series*, 148:175–194, 2003. astro-ph/0302209.
- [Starkman and Schwarz, 2005] G. D. Starkman and D. J. Schwarz. Is the universe out of tune? *Scientific American*, pages 48–55, 2005.
- [Steinhardt, 1995] P. J. Steinhardt. Cosmology confronts the cosmic microwave background. *International Journal of Modern Physics*, 10:1091–1124, 1995. astro-ph/0407329.

- [Stoeger *et al.*, 1995] W. Stoeger, R. Maartens, and G. F. R. Ellis. Proving almost-homogeneity of the universe: an almost-Ehlers, Geren and Sachs theorem. *Astrophysical Journal*, 443:1–5, 1995. gr-qc/0508100.
- [Stoeger *et al.*, 2004] W. R. Stoeger, G. F. R. Ellis, and U. Kirchner. Multiverses and cosmology: Philosophical issues. astro-ph/0407329, 2004.
- [Susskind, 2003] L. Susskind. The anthropic landscape of string theory. hep-th/0302219, 2003.
- [Susskind, 2005] L. Susskind. *The Cosmic Landscape: String Theory and the Illusion of Intelligent Design*. Little Brown, New York, 2005.
- [Susskind and Lindesay, 2004] L. Susskind and J. Lindesay. *An Introduction To Black Holes, Information And The String Theory Revolution: The Holographic Universe*. World Scientific, 2004.
- [Tegmark, 1998] M. Tegmark. Is ‘the theory of everything’ merely the ultimate ensemble theory? *Annals of Physics*, 270:1–51, 1998.
- [Tegmark, 2002] M. Tegmark. Measuring spacetime: from big bang to blackholes. *Science*, 296:1427–1433, 2002. astro-ph/0207188.
- [Tegmark, 2003] M. Tegmark. Parallel universes. In J. D. Barrow, P. C. W. D. and Harper, C., editors, *Science and Ultimate Reality: Quantum Theory, Cosmology, and Complexity*, pages 459–491. Cambridge University Press, 2003. astro-ph/0302131.
- [Tegmark *et al.*, 2004] M. Tegmark (the SDSS collaboration). Cosmological parameters from SDSS and WMAP. *Physical Review D*, 69:103501, 2004.
- [Tegmark *et al.*, 2003] M. Tegmark, A. Vilenkin and L. Pogosan. Anthropic predictions for neutrininos masses. astro-ph/0304536, 2003.
- [Tipler *et al.*, 1980] F. J. Tipler, C. J. S. Clarke, and G. F. R. Ellis. Singularities and horizons: A review article. In Held, A., editor, *General Relativity and Gravitation*. Plenum Press, 1980.
- [Tolman, 1934] R. C. Tolman. *Relativity, Thermodynamics, and Cosmology*. Clarendon Press, Oxford, 1934. Dover Publications, 1987.
- [Tryon, 1973] E. P. Tryon. Is the universe a quantum fluctuation? *Nature*, 246:396, 1973.
- [Uggla *et al.*, 2003] C. Uggla, H. van Elst, J. Wainwright, and G. F. R. Ellis. The past attractor in inhomogeneous cosmology. *Phys Rev.*, D68:103502 (1-22), 2003.
- [Uffink, 2006] J. Uffink. This vol Ch 9, 2006.
- [Vilenkin, 1982] A. Vilenkin. Creation of universes from nothing. *Physics Letters B*, 117:25, 1982.
- [Wainwright *et al.*, 1998] J. Wainwright, A. A. Coley, G. F. R. Ellis, and M. Hancock. On the isotropy of the universe: Do Bianchi VIII universes isotropize? *Classical and Quantum Gravity*, 15:331, 1998.
- [Wainwright and Ellis, 1996] J. Wainwright, and G. F. R. Ellis. *The dynamical systems approach to cosmology*. Cambridge University Press, 1996.
- [Wald, 2005] R. M. Wald. The arrow of time and the initial conditions for the universe. gr-qc/0507094, 2005.
- [Wald, 1983] R. M. Wald. Asymptotic behavior of homogeneous cosmological models in the presence of a positive cosmological constant. *Physical Review*, 28:2118, 1983.
- [Walker, 1936] A. G. Walker. On Milne’s theory of world structure. *Proceedings of the London Mathematical Society*, 42:90, 1936.
- [Walker, 1944] A. G. Walker. Completely symmetric spaces. *Journal of the London Mathematical Society*, 19:219–226, 1944.
- [Weeks *et al.*, 2003] J. Weeks, J.-P. Luminet, A. Riazuelo, and R. Lehoucq. Well-proportioned universes suppress CMB quadrupole. astro-ph/0312312, 2003.
- [Weinberg, 1989] S. W. Weinberg. The cosmological constant problem. *Reviews of Modern Physics*, 61:1–23, 1989.
- [Weinberg, 1972] S. W. Weinberg. *Gravitation and Cosmology*. Wiley, New York, 1972.
- [Weinberg, 2000a] S. W. Weinberg. The cosmological constant problem. astro-ph/0005265, 2000.
- [Weinberg, 2000b] S. W. Weinberg. A priori distribution of the cosmological constant. *Physical Review D*, 61:103505. astro-ph/0002387, 2000.
- [Wheeler, 1968] J. A. Wheeler. *Einstein’s Vision*. Springer, Berlin, 1968.
- [White *et al.*, 1993] M. White, L. M. Krauss, and J. Silk. Cosmic variance in CMB anisotropies: From 1° to COBE. *Astrophysical Journal*, 418:535–543, 1993.

- [Will, 1979] C. M. Will. The confrontation between gravitational theory and experiment. In Hawking, S. W. and Israel, W., editors, *General Relativity: an Einstein Centenary Survey*, pages 24–89. Cambridge University Press, 1979.
- [Wilson, 1989] C. Wilson. *Leibniz's Metaphysics: A Historical and Comparative Study*. Princeton University Press, 1989.
- [Wright, 2006] E. L. Wright. A century of cosmology. Astro-ph/0603750.
- [Yoshimura, 1988] M. Yoshimura. The universe as a laboratory for high energy physics. In Fang, L.-Z. and Zee, A., editors, *Cosmology and Particle Physics*, pages 293–340. Gordon and Breach, 1988.
- [Zeh, 1992] H. D. Zeh. *The Physical Basis of the direction of Time*. Springer Verlag, Berlin, 1992.
- [Zinkernagel, 2002] H. Zinkernagel. Cosmology, particles, and the unity of science. *Studies in History and Philosophy of Modern Physics*, 33:493–516, 2002.

Issues in the Philosophy of Cosmology

SUMMARY TABLE OF ISSUES AND THESES

Issue A: The uniqueness of the universe

Thesis A1: The universe itself cannot be subjected to physical experimentation

Thesis A2: The universe cannot be observationally compared with other universes

Thesis A3: The concept of 'Laws of Physics' that apply to only one object is questionable

Thesis A4: The concept of probability is problematic in the context of existence of only one object

Issue B: The large scale of the Universe in space and time

Thesis B1: Astronomical observations are confined to the past null cone, and fade with distance

Thesis B2: 'Geological' type observations can probe the distant past of our past world line

Thesis B3: Establishing a Robertson-Walker geometry relies on plausible philosophical assumptions

Thesis B4: Interpreting cosmological observations depends on astrophysical understanding

Thesis B5: A key test for cosmology is that the age of the universe must be greater than the ages of stars

Thesis B6: Horizons limit our ability to observationally determine the very large scale geometry of the universe

Thesis B7: We have made great progress towards observational completeness

Issue C: The unbound energies in the early universe

Thesis C1: The Physics Horizon limits our knowledge of physics relevant to the very early universe

Thesis C2: The unknown nature of the inflaton means inflationary universe proposals are incomplete

Issue D: Explaining the universe — the question of origins

Thesis D1: An initial singularity may or may not have occurred

Thesis D2: Testable physics cannot explain the initial state and hence specific nature of the universe

Thesis D3: The initial state of the universe may have been special or general

Issue E: The Universe as the background for existence

Thesis E1: Physical laws may depend on the nature of the universe

Thesis E2: We cannot take the nature of the laws of physics for granted

Thesis E3: Physical novelty emerges in the expanding universe

Issue F: The explicit philosophical basis

Thesis F1: Philosophical choices necessarily underly cosmological theory

Thesis F2: Criteria for choice between theories cannot be scientifically chosen or validated

Thesis F3: Conflicts will inevitably arise in applying criteria for satisfactory theories

Thesis F4: The physical reason for believing in inflation is its explanatory power re structure growth.

Thesis F5: Cosmological theory can have a wide or narrow scope of enquiry

Thesis F6: Reality is not fully reflected in either observations or theoretical models

Issue G: The Anthropic question: fine tuning for life

Thesis G1: Life is possible because both the laws of physics and initial conditions have a very special nature

Thesis G2: Metaphysical uncertainty remains about ultimate causation in cosmology

Issue H: The possible existence of multiverses

Thesis H1: The Multiverse proposal is unprovable by observation or experiment

Thesis H2: Probability-based arguments cannot demonstrate the existence of multiverses

Thesis H3: Multiverses are a philosophical rather than scientific proposal

Thesis H4: The underlying physics paradigm of cosmology could be extended to include biological insights

Issue I: The natures of existence

Thesis I1: We do not understand the dominant dynamical matter components of the universe at early or late times

Thesis I2: The often claimed physical existence of infinities is questionable

Thesis I3: A deep issue underlying the nature of cosmology is the nature of the laws of physics.

Thesis of Uncertainty: Ultimate uncertainty is one of the key aspects of cosmology

This page intentionally left blank

QUANTUM GRAVITY

Carlo Rovelli

1 INTRODUCTION

Quantum gravity (QG) is the problem of finding a theory that describes the quantum effects on gravity. These effects escape the currently accepted physical theories. Our present knowledge of the basic dynamical laws is given by quantum mechanics (QM) and quantum field theory (QFT), general relativity (GR), and the standard model of particle physics. This set of theories has obtained an empirical success nearly unique in the history of science: so far there is no evidence of observed phenomena that clearly escape or contradict this set of theories — or a minor modification of the same. But these theories become meaningless in the regimes where relativistic quantum gravitational effects are expected to become relevant. These effects are not currently observed; they are negligible at currently accessible scales and are expected to become relevant only in extreme physical regimes. For instance, they should govern the end of the evaporation of black holes, the beginning of the life of the Universe near the Big Bang, and any measurement involving an extremely short length scale ($\sim 10^{-33}$ cm, the “Planck scale”) or a very high energy. “Quantum gravity” is the name given to the theory-to-be-found that should describe these regimes.

The interest of the problem, however, goes far beyond the description of some so far inaccessible physics. The physics of the early twentieth century has modified the roots of our understanding of the physical world. It has changed the very meaning of the concepts we use to grasp it. GR, which is the field theory that describes gravity when we can disregard its quantum properties, has changed our understanding of space and time. QM, which has replaced classical mechanics as our general theory of motion, has modified the notions of matter, field, and causality. At present, we haven’t yet found a consistent conceptual frame in which these modifications make sense together. Thus, our understanding of the physical world is currently badly fragmented. In spite of its empirical effectiveness, fundamental physics is in a phase of deep conceptual confusion. The problem of QG is to combine the insights of GR and QM into a conceptual scheme in which they can coexist. It is the problem of finding a novel picture of the world capable of bringing the twentieth century scientific revolution to an end. For this reason, many consider QG to be the most important open problem in fundamental physics.

In particular, QG is an investigation on the nature of space and time. The structure and the nature of physical space are expected to change radically at

the Planck scale; and the conventional way of conceptualizing of time evolution is expected to cease to be viable at this scale. The theory is therefore likely to require us to revise the way we think of space and time.

Research in QG has developed slowly for several decades of the twentieth century, because general relativity had little impact on the rest of physics and the interest of many physicists was concentrated on the development of quantum theory and particle physics. In the last decade, the explosion of empirical confirmations and concrete astrophysical, cosmological and even technological applications of general relativity on the one hand, and the satisfactory solution of most of the particle physics puzzles in the context of the particle physics “standard model” on the other, have led to a strong concentration of interest in quantum gravity, and the progressed has become rapid. Research is currently very active.

A few tentative theories of QG have been proposed. The best developed of these are string theory [Green *et al.*, 1987] and loop quantum gravity [Rovelli, 2004]. Other very active directions of investigation include noncommutative geometry [Connes, 1994], dynamical triangulations [Ambjorn *et al.*, 1997], the spinfoam formalism [Perez, 2003] (strictly connected to the loop approach) and effective theories. Currently, none of these approaches has found any empirical corroboration, and none has won general theoretical consensus.

The problem of QG raises basic methodological issues and involves conceptual and foundational questions. Some of these are similar to the foundational questions that physics addressed at the time of other major conceptual shifts — the birth of classical mechanics, field theory, relativity, or quantum mechanics. Old problems demand new answers, in the light of the twentieth century’s novel insights. A characteristic example is a revival of the cartesian-newtonian-leibnizian debate on the relational nature of space.

No exhaustive discussion on our current understanding of the physical world, and in particular on the current knowledge about space and time, can disregard the issues and questions raised by this search.

1.1 *Quantum spacetime*

GR and QM have widely extended our understanding of the physical world. They are solidly supported empirically, and have vast scientific and technological applications. But they have destroyed the coherent picture of the world provided by prerelativistic classical physics because each of the two is formulated under assumptions contradicted by the other theory. QM is formulated using an external time variable, the t of the Schrödinger equation — or, in the case of QFT, using a fixed, nondynamical background spacetime. Both an external time variable and a fixed background spacetime are incompatible with GR.

In turn, GR is formulated in terms of Riemannian geometry: the gravitational field is assumed to be a classical deterministic dynamical field, which can be identified with Riemann’s metric field. But QM requires all dynamical fields to have

quantum properties. At small scales a field appears as made up of discrete quanta and is governed by probabilistic laws.

Thus, GR and QM are formulated in terms of mutually contradictory assumptions. In spite of their empirical success, they offer a rather schizophrenic and confused understanding of the physical world.

Roughly speaking, we learn from GR that spacetime is a dynamical field and we learn from QM that a dynamical field is quantized. Therefore at small scales we might expect a “quantum spacetime” formed by “quanta of space”, and allowing “quantum superposition of spaces”. The problem of QG is to give a mathematical and conceptual meaning to such a notion of a quantum spacetime.

Some general indications about the nature of quantum spacetime, and on the problems this notion raises, can be obtained from elementary considerations based on GR and QM. The size of quantum mechanical effects is determined by Planck’s constant \hbar . The strength of the gravitational force is determined by Newton’s constant G , and the relativistic domain is determined by the speed of light c . By combining these three fundamental constants we obtain a length, called the Planck length $l_P = \sqrt{\hbar G/c^3} \sim 10^{-33}$ cm. Quantum-gravitational effects are likely to be negligible at distances much larger than l_P , because at these scales we can neglect quantities of the order of G , \hbar or $1/c$. Therefore we expect that the GR description of spacetime as a Riemannian space holds at scales larger than l_P and breaks down approaching this scale, where the full structure of quantum spacetime becomes relevant. QG is therefore the study of the structure of spacetime at the Planck scale.

Simple arguments indicate that l_P may play the role of a *minimal* length, in the same sense in which c is the maximal velocity and \hbar the minimal exchanged action. For instance, the Heisenberg principle requires that the position of an object of mass m can only be determined with uncertainty x satisfying $mvx > \hbar$, where v is the uncertainty in the velocity; special relativity requires $v < c$; and according to GR there is a limit to the amount of mass we can concentrate in a region of size x , given by $x > Gm/c^2$, after which the region itself collapses into a black hole, removing itself from our observation. Combining these inequalities we obtain $x > l_P$; that is, gravity, relativity and quantum theory, taken together, appear to prevent position to be determined more precisely than the Planck scale. Various considerations of this kind have suggested that space might not be infinitely divisible. It may have a quantum granularity at the Planck scale, analogous to the granularity of the energy in a quantum oscillator. This granularity of space is fully realized in certain QG theories, such as loop quantum gravity, and there are hints of it also in string theory [Amati *et al.*, 1989]. Since this is a quantum granularity, it escapes the traditional objections to the atomic nature of space.

Time is affected even more radically by the quantization of gravity. In conventional QM time is treated as an external parameter and transition probabilities change in time. In GR there is no external time parameter. Coordinate time is a gauge variable which is not observable, and the physical variable measured by a clock is a complicated function of the gravitational field. Fundamental equations

of QG might therefore not be written as evolution equations in an observable time variable. Strictly speaking, this is already true in classical GR: GR does not describe evolution of physical variables in time — it describes the relative evolution of physical variables with respect to one another. But a temporal interpretation is still available in classical GR, because spacetime appears as a solution of the dynamical equations of the gravitational field. However, a solution of the dynamical equation is like a “trajectory” of a particle and in quantum theory there are no physical trajectories: there are only transition probabilities between observable eigenvalues. Therefore in QG it may be impossible to describe the world in terms of a spacetime, in the same sense in which the motion of a quantum electron cannot be described in terms of a single trajectory. It is possible that to make sense of the world at the Planck scale, and to find a consistent conceptual framework for GR and QM, we may have to give up the notion of time altogether, and learn ways to describe the world in atemporal terms. Time might be a useful concept only within an approximate description of the physical reality.

The following section sketches the historical development of QG research and illustrates the main ideas, lines of research and current tentative theories. Various issues raised by this research are then illustrated in Section 3. Section 4 discusses in particular the changes in the notions of space and time forced by QG and Section 5 discusses the relation between the problem of QG and other major open problems in fundamental physics.

2 APPROACHES

A full account of the numerous ideas and approaches towards quantum gravity is outside the scope of this article. Only a few main research lines are illustrated here. For additional references, see the bibliographical note at the end of the article.

2.1 History and directions of research

Early ideas

The fact that the gravitational field should have quantum properties, and therefore we need a theory for describing these properties, was recognized very early. Already in 1916, one year after the birth of GR, Einstein pointed out that quantum effects must lead to modifications of GR [Einstein, 1916]. In 1927 Oskar Klein suggested that QG might ultimately modify the concepts of space and time [Klein, 1927].

In the early thirties Rosenfeld [Rosenfeld, 1930b; Rosenfeld, 1930a] wrote the first technical papers on QG, soon followed by Fierz and Pauli [Fierz, 1939; Pauli and Fierz, 1939] and later Gupta [Gupta, 1952]. The idea is to introduce a fictitious “flat space”, to consider the small fluctuations of the metric around it — gravitational waves moving on flat space, described by the linearized Einstein equations — and to quantize these waves following the methods that had worked

for the electromagnetic field. More precisely, the metric field $g_{\mu\nu}(x)$, which in Einstein's theory represents at the same time the spacetime metric and the gravitational field, is written as the sum of the two terms

$$(1) \quad g_{\mu\nu}(x) = \eta_{\mu\nu}(x) + h_{\mu\nu}(x).$$

$\eta_{\mu\nu}(x)$ is interpreted as the metric of a fixed background spacetime; $h_{\mu\nu}(x)$ is interpreted as the gravitational field, and quantized. A Hilbert space of states representing quantum states of gravitational waves is introduced, where $h_{\mu\nu}(x)$ is represented by a field operator $\hat{h}_{\mu\nu}(x)$. This is called the "covariant approach" to QG. The quantum of the field $h_{\mu\nu}(x)$, which is the gravitational analog of the photon, is called the "graviton", a name already in use in the early thirties.

In 1938, Heisenberg pointed out that the fact that the gravitational coupling constant is dimensional is likely to cause problems with the quantum theory of the gravitational field [Heisenberg, 1938]. In the mid thirties a young Russian physicist, Matvei Petrovich Bronstein, realized that the unique features of gravitation require a special treatment, when the full nonlinear theory is taken into account. He realized that field quantization techniques must be generalized in such a way as to be applicable in the absence of a background spacetime. Bronstein understood that the limitation posed by GR on the mass density radically distinguishes the theory from quantum electrodynamics and would ultimately lead to the need to "reject Riemannian geometry" and perhaps also to "reject our ordinary concepts of space and time" [Bronstein, 1936].

A second line of investigation was opened in the forties by Peter Bergmann and his group [Bergmann, 1949a; Bergmann, 1949b; Bergmann, 1958; Bergmann, 1961; Bergmann and Komar, 1980]. The idea is to study and quantize the *hamiltonian* formulation of full GR, not just its linearization around flat space. This approach has the advantage that it does not assume a background spacetime on which to define the theory. The idea is that the states in the Hilbert space represent the quantum states of spacetime itself, and the full spacetime metric (maybe up to its nondynamical components) becomes a quantum operator

$$(2) \quad g_{\mu\nu}(x) \longrightarrow \hat{g}_{\mu\nu}(x).$$

This is called the "canonical approach" to QG. The same program was started independently by Dirac, who develops his constrained hamiltonian systems theory for this task [Dirac, 1950; Dirac, 1964; Dirac, 1958; Dirac, 1959].

A third approach to QG was introduced in the late fifties by Charles Misner [Misner, 1957], following a suggestion by John Wheeler. It is a quantization of general relativity à la Feynman, formally defined by the "path integral over geometries"

$$(3) \quad Z = \int Dg \, e^{-iS[g]}$$

where g is the metric field and $S[g]$ is the action of GR.

These three lines of research — covariant, canonical, and path integral — represented by equations (1), (2) and (3) respectively, still continue today. They have often influenced one another and have at times partially merged, but they have maintained a distinct flavor across more than half a century of research and are still clearly recognizable.

The basic program of the three approaches was already clearly established at the end of the fifties. The implementation of the initial programs has turned out to be a rather formidable task, but was accomplished during the sixties, in particular with the writing of the full set of Feynman rules in the covariant approach and the Wheeler-DeWitt equation in the canonical approach. Each approach, however, met serious stumbling blocks in the seventies: non-renormalizability in the covariant approach, ill-defined equations in the canonical and path integral approach. In the eighties, these stumbling blocks were overcome, in particular with the discovery of string theory in the covariant direction, and of loop quantum gravity and the (related) spinfoam formalism in the canonical and path integral directions.

The main lines of development are illustrated below.

Feynman rules and nonrenormalizability

The covariant formalism was developed during the sixties by Feynman [Feynman, 1963], DeWitt [DeWitt, 1964a; DeWitt, 1964b; DeWitt, 1965], Faddeev and Popov [Faddeev and Popov, 1967]. The technical difficulties derive from the gauge invariance of the Einstein equations, and is solved with the introduction of “ghost” particles, leading to the complete and consistent set of Feynman rules for perturbative quantum GR [DeWitt, 1967b; DeWitt, 1967c; Faddeev and Popov, 1967].

But in the early seventies the works of t’Hooft, Veltman, and then Deser and Van Nieuwenhuizen found indications that the theory does not work [t’Hooft, 1973; t’Hooft and Veltman, 1974; Deser and van Nieuwenhuizen, 1974a; Deser and van Nieuwenhuizen, 1974b], realizing Heisenberg’s early fears. The reason is that the renormalization procedure, namely the technique used in QFT to remove the infinities that appear when considering the effects of arbitrarily small (“ultraviolet”) field fluctuations, fails in the case of gravity. The definitive rigorous proof that the covariant quantization of general relativity fails because of nonrenormalizable ultraviolet divergences was obtained only later, in the late eighties, by Goroff and Sagnotti [Goroff and Sagnotti, 1985; Goroff and Sagnotti, 1986].

The interpretation of this failure is still controversial. There are two possibilities. One possibility is that the mistake was to assume, to start with, the existence of a background spacetime. The infinities come from short-distance fluctuations of the quantum field. These exist only if spacetime is continuous down to arbitrarily small scales. But the very fact that gravity is quantized questions the existence of such arbitrarily small scales. If, instead, spacetime has a quantized granular short-scale structure, the infinities might be just an artifact of the approximation taken in equation (1) considering $\eta_{\mu\nu}(x)$ (instead of the full $g_{\mu\nu}(x)$) as the spacetime metric. If so, the way out from the difficulty is to discard the background spacetime, and

quantize the full gravitational field, as is done in the canonical or path integral approaches.

The alternative possibility is that it is GR which is not the correct theory. GR is strongly empirically supported, but only at large distances. At short distances, the world might be described by a modification of GR, with better ultraviolet behavior. There is a historical precedent: Fermi theory of weak interactions was an empirically successful but non-renormalizable theory. In the case of Fermi theory, the successful solution of the problem was to replace the theory with the Glashow-Weinberg-Salam electroweak theory, which is renormalizable and corrects Fermi theory at short scales.

Motivated by this analogy, the search for a short scale correction of GR, having better finiteness properties, has spanned several decades. After numerous attempts — some of which, like supergravity [Freedman *et al.*, 1976], and high derivative theories [Stelle, 1977] raised much hope, later disappointed — the search has led to string theory.

Wheeler DeWitt theory

During the fifties and sixties, Bergmann's group and Dirac independently unraveled the hamiltonian structure of full GR, a rather complicated task. This structure was later clarified in the work of Arnowit Deser and Misner [Arnowitt *et al.*, 1962], using the metric $q_{ab}(x)$ of a constant-time spacelike surface, Ashtekar [Ashtekar, 1986; Ashtekar, 1987], using a connection field analogous to a Yang Mills field, and others.

In the early sixties, building on these results, Peres writes the Hamilton-Jacobi equations of GR [Peres, 1962]

$$(4) \quad G^2(q_{ab}q_{cd} - \frac{1}{2}q_{ac}q_{bd}) \frac{\delta S(q)}{\delta q_{ac}} \frac{\delta S(q)}{\delta q_{bd}} + \det q R = 0,$$

where R is its Ricci scalar curvature of the metric q_{ab} , and $S(q)$ is the Hamilton-Jacobi functional. In 1967, Bryce DeWitt and John Wheeler wrote the “Einstein-Schrödinger equation” [DeWitt, 1967a] following the steps taken by Schrödinger in deriving the Schrödinger equation from the Hamilton-Jacobi equation, namely interpreting the Hamilton-Jacobi equation as the eikonal approximation of a wave equation obtained replacing derivatives with derivative operators:

$$(5) \quad \left((\hbar G)^2 (q_{ab}q_{cd} - \frac{1}{2}q_{ac}q_{bd}) \frac{\delta}{\delta q_{ac}} \frac{\delta}{\delta q_{bd}} - \det q R \right) \Psi(q) = 0.$$

Today this is called the “Wheeler-DeWitt equation”. In principle, this equation is expected to describe the full quantum dynamics of gravity. In practice, the equation remained very ill-defined for a long time: until the late eighties, when Ted Jacobson and Lee Smolin [Jacobson and Smolin, 1988] find some loop-like solutions to this equation, reformulated in Ashtekar's connection formalism, thus opening the way to loop quantum gravity.

Notice that the coordinate time variable t does not appear in the classical equation (4) nor in the quantum equation (5). This disappearance of the time variable has raised an intense debate, and will be discussed below.

Misner-Hawking sum-over-geometries

In the seventies, Steven Hawking and his group [Hawking, 1979] revived and developed the Wheeler-Misner path integral (3) in the form of a “Euclidean” integral over Riemannian (namely positive definite, as opposed to pseudo-Riemannian) metrics

$$(6) \quad Z = \int Dg e^{-S[g]}.$$

The hope was that the Euclidean functional integral would prove to be a better calculational tool than the Wheeler-DeWitt equation. Hartle and Hawking [Hartle and Hawking, 1983] introduced the notion of the “wave function of the universe” and the “no-boundary” boundary condition for the Hawking integral, opening up a new intuition on QG and quantum cosmology. Jim Hartle [Hartle, 1995] developed the idea of a sum-over-histories formulation of GR into a fully fledged extension of quantum mechanics to the general covariant setting. The idea was later developed and formalized by Chris Isham [Isham, 1991].

But the Euclidean integral does not provide a way of computing genuine field theoretical quantities in QG any better than the Wheeler-DeWitt equation, and the atmosphere in QG in the middle of the eighties was rather gloomy.

The idea of a sum-over-histories definition of QG was revived again in the mid nineties by the spinfoam formalism, which offer a discretized definition of the integral (3) that appears to be better defined thanks to the same short-scale spacetime discreteness implemented in loop quantum gravity.

Black hole thermodynamics

In 1974 Hawking [Hawking, 1974; Hawking, 1975] announced a theoretical derivation of black hole radiation. A (macroscopic) Schwarzschild black hole of mass M emits thermal radiation at the temperature $T = \hbar/8\pi kGM$ (k is the Boltzmann constant). The result came as a surprise, anticipated only by the observation by Bekenstein, a year earlier, that entropy is naturally associated to a black hole [Bekenstein, 1972; Bekenstein, 1973; Bekenstein, 1974] and by the Bardeen-Carter-Hawking analysis of the analogy between laws of thermodynamics and dynamical behavior of black holes [Bardeen *et al.*, 1973]. In the light of Hawking’s result, the Bekenstein entropy of a Schwarzschild black hole is

$$(7) \quad S = \frac{kc^3 A}{4\hbar G}$$

where A is the area of the black-hole surface. Hawking’s beautiful result is not directly connected to quantum gravity — it is a skillful application of QFT in

curved spacetime, namely QFT interacting with a fixed, non quantized, gravitational field — but has a very strong impact on the field of QG. It opens a new field of research — “black hole thermodynamics” — and it opens the quantum-gravitational problems of understanding the statistical origin of the entropy (7). This is a challenge for any quantum theory of gravity.

Two years later, an influential paper by Bill Unruh [1976] convincingly argued that an observer that accelerates in the vacuum state of a conventional QFT interacts with the quantum field as if this was in a thermal bath. This shed light on black hole radiation, because an observer that remains at a fixed distance from a black hole is in constant acceleration (in order not to freely fall), and therefore black hole radiation can be interpreted simply as an Unruh effect. But at the same time this result appears to suggest that there is a deep general relation, which we do not yet understand, tying together gravity, thermodynamics and quantum theory.

In recent years, both the string and the loop approach to QG have been able to derive equation (7) from first principles [Strominger and Vafa, 1996; Rovelli, 1996a; Krasnov, 1997; Ashtekar *et al.*, 1998]. This has been considered a major success for both approaches. However, neither derivation is fully satisfactory. The string derivation does not work for conventional black holes such as a Schwarzschild black hole, but only for certain exotic black holes called extremal or nearly extremal; the loop derivation gives a finite result but the result depends on a free parameter of the theory (called the Immirzi parameter γ) that must be appropriately chosen in order to give the factor $1/4$ in (7).

Noncommutative geometry

A geometrical space M admits two alternative descriptions. One is as a set of points x , the other is in terms of a commutative algebra A of functions on M . In particular, a celebrated result by Gelfand shows that a (compact Hausdorff) space M is determined by the abstract algebra A isomorphic to the algebra of the continuous functions on M . This algebraic point of view leads to a generalization of the notion of space, obtained by considering noncommutative algebras. In this sense, a noncommutative algebra defines a “noncommutative space”.

Quantum theory is the discovery that the phase space of a dynamical system (the set formed by its classical states) must be replaced by a noncommutative space. In fact, the system’s observables — that represent the ways we can interact with the system — form a commutative algebra of functions on the classical phase space, which becomes a noncommutative algebra in QM.

In the case of physical space, A can be identified with an algebra of coordinates, or with momentum space. If we interpret the elements of A as representing physical measurements, it is natural, in the light of quantum theory, to consider the possibility that the algebra be noncommutative. Accordingly, the hypothesis has been made that the short-scale structure of physical space might be described by a noncommutative geometry. This idea has been explored in a number of variants [Doplicher *et al.*, 1994; Doplicher *et al.*, 1995;

Doplicher, 1996].

A connection with GR has appeared in the approach developed by Alain Connes [Connes, 1994]. Connes has noticed that in the algebraic framework the notion of distance is naturally encoded in the Dirac operator D . This is the derivative operator that appears in Dirac's spinor field equation for an electron. Let \mathcal{H} be the Hilbert space formed by the spinor fields on a given Riemannian (spin) manifold M , D be the (curved) Dirac operator, and A an algebra of functions on M , seen as (multiplicative) operators on \mathcal{H} . From the triple (\mathcal{H}, A, D) , called a "spectral triple", we can reconstruct the Riemannian manifold. In particular, the distance between two points x and y can be obtained as

$$(8) \quad d(x, y) = \sup_{\{f \in A, \|[D, f]\| < 1\}} |x(f) - y(f)|$$

a beautiful and surprising algebraic definition of distance. A non-commutative spacetime might be described by a spectral triple in which A is non-commutative. Connes suggests that this algebra may be chosen on the basis of the symmetries of the standard model, following the idea that the standard model might reveal the short-scale structure of spacetime in the same manner in which Maxwell theory revealed the structure of Minkowskian spacetime. The Connes-Chamseddine "spectral action", is simply the trace of the Dirac operator D , $S = \text{Tr}[f(D^2/(\hbar G))]$, where f is the characteristic function of the $[0, 1]$ interval. Remarkably, this action turns out to include the standard model action, including the poorly understood Higgs sector, as well as the action of GR [Chamseddine and Connes, 1996; Chamseddine and Connes, 1997]. The precise relation between the noncommutativity of noncommutative geometry and of QM has not yet been extensively investigated.

Other ideas and directions

A large number of other ideas and directions of investigation about QG have been proposed. Some of these research directions are still active. Only a few are mentioned below.

A project extensively explored is to define a quantum gravity in terms of the continuum limit of a discrete lattice theory, a technique that works in the case of quantum chromodynamics. Various attempts in this direction have failed in the past, because the lattice theory considered turned out not to have a continuum limit. One of the versions of this program, called *dynamical triangulation* is still very active, although no proof of the existence of a continuum limit exists yet.

Raphael Sorkin and his group have long explored a discrete model in which spacetime is replaced with a discrete set of points equipped with an ordering representing the causal relations [Sorkin, 1983]. Remarkably, the model has predicted a small but non-vanishing cosmological constant, of the correct order of magnitude, a prediction recently confirmed.

Roger Penrose and his group have developed twistor theory as a reformulation of metric geometry, with the hope of addressing the QG problem [Penrose, 1967]. So far, the results of twistor theory are more of mathematical than physical relevance.

Other research directions include Hartle's quantum mechanics of spacetime [Hartle, 1995], quantum Regge calculus [Williams and Tuckey, 1992; Williams, 1997], 't Hooft's deterministic approach ['t Hooft, 1996] and Finkelstein's theory [Finkelstein, 1997].

“Phenomenology” and Lorentz invariance

Until a few years ago, the research community was convinced that QG effects were certainly far outside our current observational reach. This conviction has been shaken by a number of recent suggestions that these effects might in fact be on the verge of being observable. The suggestion has even been made that certain data already observed, and which appear to be difficult to interpret with conventional physics, might be affected by QG effects. These suggestions concern for instance the cosmic propagation of high energy particles, fine details in the cosmological density spectrum, and others.

The issue appears to be related to the problem whether QG breaks Lorentz invariance. Small Lorentz noninvariant QG effects, if they exist, could be within or near observational reach. Lorentz invariant effects, on the other hand, are presumably far smaller, because Lorentz invariance forbids certain effects to happen. For instance, a small deviation from the Lorentz invariant dispersion relation governing light propagation could accumulate over cosmological travel times and yield observable frequency-dependent delays. In a Lorentz invariant context, light travel-time is a meaningless notion.

Naively one might expect that the existence of a minimal length in QG necessarily breaks Lorentz invariance. The argument is that the minimal length must be Lorentz contracted under a change of inertial frame, and therefore could not be minimal. But this argument is incorrect because it disregards quantum theory [Rovelli and Speziale, 2003]. In quantum theory a discrete quantity appears as the *eigenvalue* of an observable quantity, while a symmetry transformation transforms states, and therefore *means values*, not eigenvalues.

To illustrate this phenomenon, recall that in classical mechanics the z component L_z of the angular momentum transforms continuously under rotations. In the quantum theory, let a system be in the eigenstate $|\psi\rangle = |\hbar/2\rangle$ of L_z . Seen from a rotated reference frame, this system will appear to be in a superposition $|\psi\rangle = \alpha|\hbar/2\rangle + \beta|-\hbar/2\rangle$, where α and β vary continuously with the rotation angle. Therefore the expectation value $\overline{L_z} = |\alpha|^2\hbar/2 - |\beta|^2\hbar/2$ varies continuously in the rotation, but the eigenvalues remain the same. In physical terms: we always observe the discrete values $L_z = \pm\hbar/2$ in all reference systems — what changes continuously in a rotation is the probability of seeing one or the other. In the same fashion, in loop quantum gravity a (nonvanishing) minimal area is an eigenvalue. A surface which is in an eigenstate of the area will appear in a superposition of different area eigenstates if seeing from a boosted reference frame. The expectation value of the area of a surface can be smaller than the minimal area, but a (nonvanishing) measurement outcome cannot.

2.2 *The main current tentative theories*

The two currently most developed and most studied quantum theories of gravity are string theory and loop quantum gravity.

Strings

The major reason for the interest in string theory is that it is a fundamental theory of the world, including the gravitational field, which is likely to be free of ultraviolet divergences, and which encodes in a natural and strictly unified structure all the diverse ingredients we find in the world.

The starting point of the theory is the hypothesis that elementary objects are not point-like particles but rather strings, namely one-dimensional objects. The theory was initially studied as a tentative theory of the strong interactions, where it turned out to be incorrect. Quantum string theory is only consistent if spacetime has a certain dimension, called the critical dimension, which is 26 for the bosonic string and 10 for the supersymmetric string which includes fermions. The problem of reconciling the critical dimension with the fact that our world appears to be four dimensional is still open.

In 1984, Green and Schwarz introduced the idea that string theory might be a unified theory of all interactions, including gravity [Green and Schwarz, 1984]. In fact, one of the vibration modes of the proposed string has spin two, and can be identified with the graviton. Furthermore, a necessary (in general not sufficient) condition for string theory to be well-defined is that the background spacetime satisfies an equation that reduces to the Einstein equation in the large distance limit.

Consistency restricts the string models to a few alternatives. A supersymmetric model defined on a 10-dimensional flat spacetime using a large gauge group, appears to include all the ingredients of our world: the gauge group includes a subgroup which is the gauge group of the particle physics standard model, and the lowest energy vibration modes of the string include fermions, gauge bosons, and the graviton. Although no complete proof is available, the theory appears to have no ultraviolet divergences.

The idea is that six of the ten dimensions of spacetime may be invisible to us, because they are wrapped (“compactified”) into a very small space (or because we are constrained to live on a four-dimensional surface). The effective physical theory in the four visible dimensions depends on the way the six extra dimensions are compactified. This can happen in a great number of different manners, giving rise to a huge number of effective four-dimensional theories. For the moment, no selection principle among this large number of possibilities has been found. Some of the resulting low energy models appear to have a strong resemblance to the standard model, but so far none seems to give precisely the physics we observe at low energy.

String theory is defined in terms of a perturbation expansion on a 10 dimensional fixed spacetime background. In the mid nineties, several nonperturbative aspects

of string theory began to be investigated. Higher dimensional excitations, called “branes” (from “membrane”) [Polchinski, 1995] appear to be needed in the theory for consistency, besides the strings themselves. (It has been suggested that the four-dimensional surface on which we live could be a four-dimensional brane.) The different string models appear to be related to one another (and to 11-dimensional supergravity) via simple transformations called “dualities”, suggesting that all the different string models are actually different limits of a single unknown fundamental theory, tentatively-called “M-theory”. The actual construction of this hypothetical fundamental theory — expected to be background independent — is still missing, and so far string theory exists only in the form of a number of (loosely) related models defined in terms of expansions over assigned background spacetimes.

In 1998, a certain conformal field theory was shown to include a sector that appears to be related to a supergravity theory on the product of Anti-deSitter spacetime and spheres. This led to the conjecture that the compactifications of string theory on an Anti-deSitter spacetimes is “dual” to a field theory on the spacetime boundary. In turn, this led to a new proposal for defining M-theory itself in term of a boundary theory: the idea is to reach background independence for M-theory using background dependent methods for the boundary theory.

The difficulties of the theory are many. No selection mechanism for the compactification is known — this is the problem of the selection of the “vacuum”; since each compactification gives different physical predictions, and there are hundreds of thousands of possible compactifications, string theory is effectively a collection of a huge number of different theories, each with different predictions and each with different physical parameters. As a result, the theory is incapable of computing the values of the standard model parameters and almost completely nonpredictive, in the sense that it can be compatible with almost any future experimental outcome. According to some critics, this lack of predictivity undermines the very nature of string theory as a scientific theory.

Even if we are willing to choose a compactification ad hoc, no compactification giving precisely the standard model in the low-energy limit is known. The theory requires supersymmetry, and the existence of observable supersymmetric particles has repeatedly been claimed as the distinctive prediction of the theory; but, in spite of several preliminary announcements, supersymmetric particles have not been found in experimental particle physics. Similarly, the possibility of detecting effects of the invisible dimensions has been considered, but experiments have given negative results. The theory requires a huge baggage of new physics (extra dimensions, an infinite number of fields with arbitrary masses and spins, supersymmetric particles ...) but so far none of this appears to be present, or have observable consequences, in the real world.

Loops

The main reasons for interest in loop quantum gravity are: that its physical assumptions are only QM and GR, namely well-tested theories; the fact that the theory is background independent; and that it is a well developed attempt to incorporate the general relativistic notions of space and time into QFT. The theory makes no claim of being a final “Theory Of Everything”. It is ultraviolet finite, without requiring high-energy modifications of GR, supersymmetry, extra dimensions, or other unobserved physics.

Loop quantum gravity was introduced in 1988. The theory is the result of the merging of two lines of research, which turn out to solve each others difficulties [Rovelli and Smolin, 1988; Rovelli and Smolin, 1990].

The first of these was the Wheeler-deWitt theory. As in the Wheeler-deWitt approach, loop quantum gravity is a straightforward quantization of GR, with its conventional matter couplings, and is based on no specific physical assumption other than GR and QM. Following the basic rules of QM, the quantum states of loop quantum gravity are obtained from a representation of an algebra of field variables of GR; their physical interpretation is obtained by diagonalizing self-adjoint operators that represent physical quantities. The difference with respect to the old Wheeler-DeWitt theory is in the choice of an algebra of loop-variables as basic variables for the quantization. Thanks to this, the ill-defined Wheeler-DeWitt theory becomes a well-defined formalism where finite physical quantities can be computed.

The second input was the idea that gauge theories are naturally described in terms of loop-like excitations. This idea can be traced back to the very beginning of field theory, an intuition of Faraday’s. Faraday understood electric and magnetic phenomena in terms of lines, the “Faraday lines”, that fill up space. In the presence of charges, the Faraday lines can start and end on the charges; in the absence of charges, they close, forming “loops”. Maxwell translated Faraday’s intuition into mathematical physics, introducing the electric and magnetic field, which are vector fields everywhere tangent to the Faraday lines, thus opening the way to modern physics, which is entirely based on the notion of field. The idea that gauge field theories are better understood in terms of loops has been defended by many scientists, including Polyakov, Mandelstam, Wilson, and others. A quantum excitation of a single Faraday line is called a “loop state”.

A formulation of a QFT in terms of loop states is viable and well understood in the context of the lattice approximation; but it faces difficulties when defined over a continuum spacetime background. However, these difficulties disappear in a background independent context. The reason is that in the presence of a background, the loop states are localized on the background spacetime: there is a distinct state for each position of the loop in space. In the case of gravity, instead, there is no background spacetime. The loop states themselves are the quantum excitations of space. Therefore loop states are not immersed in space: they “weave-up” physical space themselves, in the same manner in which an ensemble of threads

can weave the fabric of a T-shirt.

More precisely, the loop states of QG have self-intersection points called “nodes”. A node represents an elementary quantum excitation of space, or a single atom of space. Two nodes directly connected along a loop represent adjacent atoms of space. Nodes and links connecting nodes form a graph and carry quantum numbers. These quantum numbers determine the quantized volume of the atoms of space and the quantized area of the elementary surfaces separating adjacent nodes. A graph with these quantum numbers is called a “spin-network”, because the quantum numbers on the links turn out to be half-integers, or spins.

A spin network state does not have a position. Only combinatorial relations defining the graph are significant, not its shape or its position in space. In fact, a spin network state is not *in* space: it *is* space. Hence, in spite of its conservative basic assumptions (QM and GR), loop quantum gravity leads to a radically novel picture of space.

The possible values that the volume of a physical region or the area of a physical surface can take are determined by the spectra of the corresponding operators, following standard QM rules. These turn out to be discrete, giving the Planck-scale granular structure of space. These spectra have been computed and represent quantitative physical predictions of loop quantum gravity: a Planck-scale precision measurement of any area or volume is predicted to give as a result only the values in these spectra. For instance, the (main sequence of the) spectrum of the area is given by the expression [Rovelli and Smolin, 1995]

$$(9) \quad A = 8\pi\gamma\hbar G \sum_i \sqrt{j_i(j_i + 1)},$$

where j_i is an n -tuple of half-integers (corresponding to the quantum numbers of the links of the spin network state crossing the surface whose area is measured). γ is the Immirzi parameter, mentioned in Section 2.1.

The dynamics is determined by a Wheeler-deWitt equation on the space of spin network states. Its ultraviolet finiteness is a consequence of the granular structure of space. Different finite and well-defined versions of this equation have been constructed. At present it is not yet clear which of these, if any, is the physically correct one.

Applications of the theory include a derivation of the Bekenstein black hole entropy mentioned in Section 2.1, applications to the description of the classical singularities, such as the ones at the center of black holes, and applications to cosmology. The theory appears to be capable of controlling the black hole singularities and the initial Big Bang singularity. Indirect empirical evidence supporting predictions of the theory is actively searched in the astrophysical and cosmological domains.

The main difficulties of loop quantum gravity lie in recovering low energy phenomenology. Quantum states corresponding to the Minkowski vacuum and its excitation have not yet been constructed, and particle scattering amplitudes have not been computed. This deficiency weakens the strength of the finiteness claim,

and bears on one of the key requirements on a quantum theory of gravity: full recovery of low energy physics. The dynamics is still poorly understood: the Wheeler-deWitt equation exists in more than one version. The lack of unitary evolution in time and the overall radical conceptual novelty of the results of the theory, where background spacetime is discarded altogether, are questioned by some.

The loop-string debate

A theory begins to be credible only when its original predictions are reasonably unique and are confirmed by new experiments. Neither loop quantum gravity nor string theory — nor any other current tentative theory of QG — are yet credible in this sense. Furthermore, in spite of much effort, both theories are still badly incomplete and far from being clearly understood. The problem of QG must therefore be considered still fully open.

Nevertheless, in both directions the research has progressed considerably in recent years: many problems that appeared too hard ten years ago have now been solved, and incomplete but *possible* solutions of the QG puzzle are now at hand.

However, the two theories differ profoundly in their hypotheses, achievements, specific results, and in the conceptual frame they propose. The issues they raise concern the foundations of the physical picture of the world, and the debate between the two approaches involves conceptual, methodological and philosophical issues.

The lesson of string theory appears to be that in order to remove the difficulties of the perturbative quantization of GR we have to couple the gravitational field to matter. Finiteness is achieved by replacing the pointlike Feynman vertices of conventional QFT with non-point-like interactions between strings, which are extended objects. The theory preserves the basic conceptual structure of QFT (background spacetime, unitarity, predictions in terms of an asymptotic S-matrix. . .) at the prices of renouncing a full implementation of the general covariance that characterizes GR, of huge extra baggage (extra dimensions, supersymmetry, infinite fields. . .) and of a dramatic decrease in predictiveness.

Loop quantum gravity, on the other hand, is rooted in the general covariance that characterizes GR. Ultraviolet finiteness is a consequence of the granular structure of space, which, in turn, is a standard quantum mechanical effect appearing when we regard GR as a theory of spacetime itself, and not as a theory of small perturbations around a background spacetime. The interest of the loop theory, therefore, is that it is a determined effort towards a genuine merger of QFT with the world view that we have discovered with GR. Furthermore, it leads to well-defined physical predictions which are in principle falsifiable. However, even disregarding the incompleteness of the theory, the conceptual price for this result is heavy: the theory gives up unitarity, time evolution, Poincaré invariance at the fundamental level, and the very notion that physical objects are localized in space and evolve in time.

Whether these radical conceptual steps are viable, and, if viable, whether they are justified, is a hotly debated issue.

3 METHODOLOGICAL ISSUES

3.1 Justification of the quantum gravity search

Absence of empirical data

The first obvious question about the search towards QG is whether the search is legitimate at all, given the total absence of empirical data directly about the regimes QG is concerned with. We have no direct empirical guidance in searching for QG — as, say, atomic spectra guided the discovery of quantum theory.

Some critics have argued that the QG search is futile, because anything might happen in QG regimes, at scales far removed from our experience. Maybe the search is impossible because the space of possible theories is too large.

At present, this worry is probably unjustified. If this were the problem, we would have plenty of complete, predictive and coherent theories of QG, and the problem would be the choice among them. Instead, the situation is the opposite: we haven't any. The fact is that we do have plenty of information about QG, because we have QM and we have GR. Consistency with QM and GR, plus internal consistency, form an extremely strict set of constraints. The problem currently debated is to find at least one complete and consistent theory of QG. If more will be found, we will have of course to resort to experiments to select the physically correct one.

Should gravity be quantized?

The possibility that quantum gravitational effects do not exist and gravity is intrinsically classical (non quantum) has been often suggested. The justification for this suggestion is that gravity can be seen as an interaction profoundly different from the others, since it admits a description in terms of spacetime geometry. This suggestion has also generated a research program, aiming at testing the consistency of a theory in which classical gravity interacts with QFT.

In its simpler form, this suggestion has today been largely abandoned. The reason is that, as was noticed in the early days of QM, an interaction between a classical and a quantum variable is always inconsistent. If Heisenberg uncertainty relations are violated for one dynamical variable, they are violated for all other variables as well. The idea of circumventing gravity quantization, however, has reappeared under various forms.

One suggestion is that the gravitational field may not represent true microscopic degrees of freedom, but only a collective, or “hydrodynamical”, large scale description of these. This hypothesis is supported by phenomena such as the relations between gravity and thermodynamics revealed by the Unruh effect. Ted Jacobson [Jacobson, 1995] has even been able to derive the Einstein equations from (7) and standard thermodynamical relations, providing evidence that could be interpreted

as supporting this idea. However, even if the gravitational field is just a collective variable, this does not mean that it will not display quantum effects. QM does not govern just elementary degrees of freedom; it governs all degrees of freedom, including collective ones. Thus, this possibility, even if realized, would not refute the need of a quantum theory of gravity.

Another suggestion is that gravity may be an emergent phenomenon induced by the other quantum fields. This idea is suggested by the fact that the renormalization process for a QFT on a curved spacetime generates terms in the action which are proportional to polynomials in the Riemann curvature, and the lowest order term is precisely the action of GR. The difficulty about this suggestion is that it is ambiguous as regards the dynamical status of the metric field. The variational principle states that dynamics is determined by the variation of the action with respect to the dynamical variables only, not with respect to anything appearing in it. If the metric field is assumed to be a dynamical variable, then it is a dynamical field like any other, and the fact that the dependence of the action with respect to it is modified by the renormalization of its interaction with other fields may change the details of its dynamics, but not the fact that it is a quantum field. If, on the other hand, the metric is not a dynamical field, then the action must not be varied with respect to it (as it is not varied with respect to it in the special relativistic context), and therefore the Einstein equations are not generated by the new terms in the action. In the first case the gravitational field needs to be quantized; while the second case is in contradiction with the empirical fact that the classical Einstein equations are satisfied

3.2 Research attitudes

Different attitudes can be distinguished in the physics community with respect to the methodology used for searching for a QG theory.

(a) The “pessimistic” attitude, already mentioned above, is that of those who worry that too many possibilities are open, anything might happen between here and the Planck scale, and the search for a quantum theory of gravity is therefore futile.

As mentioned, this worry is unfounded, because we do not have too many complete QG theories: we haven’t any.

(b) The view is often expressed that some totally new, radical and wild hypothesis is needed for QG. This “wild” attitude is based on the observation that great scientists had the courage to break with old and respected assumptions and to explore some novel “strange” hypotheses. From this observation, the “wild” scientist concludes that any strange hypothesis deserves to be investigated, even if it violates well established facts.

On historical grounds, this expectation is probably ill-founded. Wild ideas pulled out of the air have rarely made science advance. The radical hypotheses that physics has successfully adopted have always been reluctantly adopted because they were forced by new empirical data — Kepler’s ellipses, Bohr’s quantiza-

tion, Planck's energy quanta — or by stringent theoretical deductions — Maxwell displacement current, Einstein's relativity. Generally, arbitrary novel hypotheses have led nowhere. This consideration leads to the next attitude in (c).

(c) Part of the research in QG is motivated by the hope that the knowledge of the world coded into GR and QM can be a good guide for finding a theory capable of describing physical regimes that we have not yet explored.

A motivation for this hope is that today we are precisely in one of the typical situations in which theoretical physics has worked at its best in the past. Many of the most striking advances in theoretical physics have derived from the effort to find a common theoretical framework for two basic and apparently conflicting discoveries. For instance, the aim of combining special relativity and non-relativistic quantum theory led to the theoretical discovery of antiparticles; combining special relativity with Newtonian gravity led to general relativity; combining the Keplerian orbits with Galilean physics led to Newton's mechanics; combining Maxwell theory with Galilean relativity led to special relativity, and so on. In all these cases, major advances have been obtained by "taking seriously" apparently conflicting theories, and exploring the implications of holding true the essential tenets of both theories. Today we are in one of these characteristic situations. We have learned two new very general "facts" about nature, expressed by QM and GR: we have "just" to figure out what they imply, taken together.

(d) A different point of view on the problem is held by those who accept that QM has been a conceptual revolution, but do not view GR in the same way. According to this point of view, the discovery of GR was "just" the writing of one more classical field theory. This field theory is likely to be only an approximation to a theory we do not yet know, and its teachings should not be overestimated. According to this opinion, GR should not be taken too seriously as a guidance for theoretical developments.

A possible objection to this point of view is that it derives from the confusion between (i) the specific form of the GR action and the GR field equations and (ii) the modification of the notions of space and time engendered by GR. The GR action could be a low energy approximation of something else. But the modification of the notions of space and time has to do with the diffeomorphism invariance and the background independence of the theory, not with its specific form. The challenge of QG is to incorporate this novelty into QFT, not the specific form of the GR action.

(e) A common attitude is the "pragmatic" attitude of the physicist who prefers to disregard or postpone these foundational issues and, instead, develop and adjust current theories. This style of research was effective during the sixties in the search for the particle physics standard model, where a long process of adjustment of existing QFT's led to a very effective theory.

It is questionable whether this attitude could be effective in a situation of foundational confusion like the present one. During the sixties empirical data were flowing in daily, to keep research on track. Today no new data are available. The "pragmatic" attitude may mislead the research: in the extreme case, the "prag-

matic” physicist focuses only on the development of the theory at hand, without caring if the world predicted by the theory resembles less and less the world we see. Sometimes he is even excited that the theory looks so different from the world, thinking that this is evidence of how far ahead he has advanced in knowledge. But it is more likely that the difference between the theory and the world is only evidence of how much he is lost. Unfortunately similar excesses plague theoretical research today.

The cumulative aspect of scientific knowledge and influence of the philosophy of science

The “pessimistic”, “wild” and “pragmatic” attitudes illustrated above may have been influenced by a philosophy of science that under-emphasizes the cumulative aspect of scientific knowledge, and emphasizes, instead the “incommensurability” between an old theory and a new theory that historically supersedes it. More or less informed awareness of this long standing debate in philosophy of science has indeed affected the research attitude of many theoreticians.

On the other hand, attitude (c) described above is based on the expectation that the central physical tenets of QM and GR represent our best guide for accessing the unexplored territories of the quantum-gravitational regime. In a difficult research situation where cataclysmatic evolution is expected anyway in the consequences of the theory (for instance, the change of the nature of space), conservative assumptions based on the confidence on the cumulative aspect of knowledge can play an important role.

This faith in a cumulative aspect of scientific knowledge is based on the idea that there are discoveries that are “forever”. For instance, the Earth is not the center of the universe, simultaneity is relative, absolute velocity is meaningless, and we do not get rain by dancing.

The fact that major aspects of a theory can have value outside the domain for which the theory was discovered may be at the root of much of the historical effectiveness of theoretical physics, and in particular of spectacular predictions such as Maxwell’s radio waves, Dirac’s antimatter or GR’s black holes. This can perhaps be understood as just scientific induction: as a consequence of the fact that Nature has regularities. This is not the place to enter this discussion; but it is relevant to remark that the existence of these regularities is held by several researchers in QG as a source of confidence — although, of course, not certainty — that the basic facts about the world found with QM and GR will be confirmed, not violated, in the quantum gravitational regimes that we have not yet empirically probed.

4 THE NATURE OF SPACE AND TIME

GR has modified the way we understand space and time. Combining GR with QM requires a further modification of these notions. It is important, however, to

clearly distinguish the modifications of the notions of space and time required by QG from the ones already implied by GR alone. These are briefly summarized in Section 4.1 below. Section 4.2 and 4.3 then discuss the notions of space and time in QG.

4.1 *The physical meaning of GR*

GR is the discovery that spacetime and the gravitational field are the same entity. What we call “spacetime” is itself a physical object, in many respects similar to the electromagnetic field. We can say that GR is the discovery that there is no spacetime at all. What Newton called “space”, and Minkowski called “spacetime”, is nothing but a dynamical object — the gravitational field — in a regime in which we neglect its dynamics.

In newtonian and special relativistic physics, if we take away the dynamical entities — particles and fields — what remains is space and time. In general relativistic physics, if we take away the dynamical entities, nothing remains. The space and time of Newton and Minkowski are reinterpreted as a configuration of one of the fields, the gravitational field. This implies that physical entities — particles and fields — are not all immersed in space, and moving in time. They do not live on spacetime. They live, so to speak, on one another.

In classical GR it is customary to maintain the expressions “space” and “time” to indicate aspects of the gravitational field. But in the quantum theory, where the field can have quantized “granular” properties and its dynamics is quantized and therefore only probabilistic, most of the “spatial” and “temporal” features of the gravitational field are probably lost.

This absence of the familiar spacetime “stage” is called the *background independence* of the classical theory. Technically, background independence is realized by the gauge invariance of the GR action under (active) diffeomorphism. A diffeomorphism is a transformation that smoothly drags all dynamical fields and particles on the four-dimensional coordinate manifold. In turn, gauge invariance under diffeomorphism (or *diffeomorphism invariance*) is the consequence of the combination of two properties of the action: its invariance under arbitrary changes of coordinates (or *general covariance*) and the fact that there is no non-dynamical “background” field. Thus: *background independence* = *diffeomorphism invariance* = (*general covariance*+absence of non-dynamical background fields). These notions are illustrated in more detail below.

Diffeomorphism invariance

Pre-general-relativistic field theories are formulated in terms of a spacetime manifold M , and a set of fields $\varphi_1, \dots, \varphi_n$ on M . The manifold M is a (pseudo-)metric space whose points $P \in M$ represent the physical points of spacetime. Spacetime points are labelled by coordinates $x = (x^1, x^2, x^3, x^0)$ that represent the reading of measuring devices: clocks and distance-measuring devices (“rods”). More precisely, M is equipped with a (pseudo-)distance function $d(x, y)$ interpreted as

the 4-interval between the two points x and y : a negative $d^2(x, y)$ gives the time measured by a clock in inertial motion between x and y ; a positive $d^2(x, y)$ gives the proper length of a rod with the ends on x and y , in a state of inertial motion with respect to which x and y are simultaneous according to Einstein's definition of simultaneity; a null $d(x, y)$ indicates that light travels in vacuum from x to y . Notice that pre-general-relativistic physics deals with (relations between) two distinct types of measurements: (i) spacetime measurements measuring spacetime observables, performed by means of clocks and distance-measuring devices, and (ii) field measurements, measuring field observables, namely the values (or functions) of the fields $\varphi_1, \dots, \varphi_n$.

This same interpretation framework is used in special-relativistic QFT. The only difference is that field observables can be quantized. The number of excited quanta has a particle interpretation. In a typical high energy scattering experiment, for instance, the field observable (ii) is the number of particles revealed by a particle detector (which is a field measuring device); while the spacetime observable (i) is the momentum of the particle, determined by measuring the spacetime position of the detector.

The theory does not predict the value of field observables φ alone, or spacetime observables x alone, but only combinations of the two, such as the value $\varphi(x)$ of a field φ at a certain spacetime location x . Spacetime and field observables are both quantities that have a direct operational interpretation; they can be called "partial observables". On the other hand, the quantities that can be predicted by the theory, such as $\varphi(x)$, for a given position x , can be called "complete observables".

The interpretation of a *general* relativistic field theory is different. In such a theory there is a field g representing the gravitational field and possibly other fields, representing other dynamical variables. These fields are defined on a differentiable manifold M , coordinatized by coordinates x . The formal structure of a general relativistic field theory is therefore similar to the structure of a pre-general-relativistic field theory. But two major differences force a different interpretation. First, the manifold M on which the fields are defined is not a metric manifold. The gravitational field g equips M with a metric structure $d_g(x, y)$.¹ Therefore clocks and distance-measuring-devices measure properties of the gravitational field g . It follows that the distinction between spacetime observables of the kind (i) and field observables of the kind (ii) is blurred. This blurring of the distinction between the two kind of partial observables is a crucial conceptual novelty of GR.

Second, the field equations are invariant under a transformation of the fields called active diffeomorphisms. An active diffeomorphism $g \rightarrow \tilde{g}$ is determined by (but should not be confused with) a smooth invertible function $f: M \rightarrow M$. Under a diffeomorphism transformation, the field g and all other fields are "dragged along" M by f . For instance, the transformed field \tilde{g} defines a new distance

¹The length of a curve $\gamma^\mu(s)$ in M is $d_g(\gamma) = \int ds \sqrt{g_{\mu\nu}(\gamma(s)) \frac{d\gamma^\mu(s)}{ds} \frac{d\gamma^\nu(s)}{ds}}$ and $d_g(x, y)$ is a local extremum of $d_g(\gamma)$ over the curves that join x and y .

function $d_{\tilde{g}}(x, y)$ which is related to the one defined by g by

$$(10) \quad d_{\tilde{g}}(f(x), f(y)) = d_g(x, y).$$

In words: the distance between the two points $f(x)$ and $f(y)$, defined by the field \tilde{g} is the same as the distance between the two points x and y defined by g .

The importance of this invariance of the field equations is due to the following. An active diffeomorphism may modify a solution of the GR equations in the future of a certain time surface t_0 , without modifying it at all in the past of t_0 . Therefore two distinct solutions of the equations of motion can be equal in the past and differ in the future. This fact gives us a choice: either (i) we interpret the theory as an indeterministic theory, where the future is not determined by the past. Or (ii) we interpret active diffeomorphisms as a “gauge invariance”: that is, we postulate that the complete observables of the theory are only given by quantities that are invariant under this transformation. The alternative (i) is not viable, because experience shows that classical gravitational physics is completely deterministic. We are therefore forced to alternative (ii), which has heavy interpretative consequences.

To understand these consequences, let P be a point of M . Let $\varphi(P)$ be any property of the fields at P . For instance, $\varphi(P)$ may represent a value of the electromagnetic field at P or the spacetime scalar curvature at P , or something similar. None of these properties is invariant under active diffeomorphism. Therefore, it follows from the argument above that none of these properties can be predicted by the theory. Therefore the theory does not determine the physics at spacetime points $P \in M$.

At first this conclusion might sound bewildering: if physics does not predict what happens at spacetime points, what can it predict? In fact, historically, Einstein himself got at first confused and frustrated by this observation, to the point of stepping back from the diffeomorphism invariance he previously expected GR to have [Norton, 1984]. Einstein’s version of the argument given above is called the “hole argument” (because Einstein considered a diffeomorphism affecting only a finite region of spacetime, empty of matter, or a “hole”), and was presented in [Einstein and Grossmann, 1914]. On this argument, and the discussion it has raised, see for instance [Earman, 1987; Earman, 1989; Belot, 1998; Earman, 2001; Pauri and Vallisneri, 2002] and references therein. Later however, Einstein changed his mind and accepted both diffeomorphism invariance and the conclusion (ii), realizing that this conclusion fully implemented his intuition on the very central physical meaning of the general relativistic conceptual revolution.

The way out of the puzzle is to understand that in the general-relativistic context the points of the manifold do not represent physical entities with an existence independent from the fields. Asking what are the properties of the fields at P is meaningless. Spacetime locations can only be determined by the fields themselves, or, by any other dynamical object we are considering. For instance, if the theory we consider includes two particles and the trajectories of the two particles happen to meet once, then the meeting of the particles determines a spacetime

point. The theory is able to predict the value of the fields and any other physical properties at the spacetime point determined by the meeting of the particles. However, this point cannot be naively identified as a point of M , because the same physical situation can be represented by the set of dynamical variables obtained by an active diffeomorphism, where the particles meet now at a different point, say Q of M . The value of the fields at the point where the particles meet is invariant under such a transformation, because the particles' trajectories and the fields are dragged along M together. Einstein called this way of determining location in terms of the dynamical objects (fields and particles) of the theory itself, "spacetime coincidences".

Thus, a general relativistic theory does not deal with values of dynamical quantities at given spacetime points: it deals with values of dynamical quantities at "where"'s and "when"'s determined by other dynamical quantities.

Strictly speaking also in the pre-general-relativistic context physics deals with values of dynamical quantities at "where"'s and "when"'s determined by other physical quantities, because the times and distances used to determine location are physical quantities. But in the pre-general-relativistic context we can make a strict separation between: (i) "spacetime", viewed as a background entity, and measured by clocks and rods that one considered non-dynamical, and (ii) dynamical variables. In the general-relativistic context, on the other hand, this separation is lost, time and distance measurements are reinterpreted as measurements of the gravitational field, on the same footing as other field measurements, and there is no distinction between non-dynamical background and dynamical physical variables.

Physical meaning of the coordinates

A consequence of the above is that in the general-relativistic context the physical interpretation of the coordinates x is different from their interpretation in the non-general-relativistic context. In the general-relativistic context the coordinates x have no interpretation at all: observable quantities in GR correspond to quantities of the theory that are independent of the coordinates x . Recall that the non-general-relativistic coordinates represent the reading of clocks and rods: in the general-relativistic context, the reading of clocks and rods is represented by the non-local function $d_g(x, y)$ of the gravitational field. The fact that the non-general-relativistic coordinates x and the general-relativistic coordinates x are denoted in the same manner is only an unfortunate historical accident.

To illustrate in which sense observable quantities are independent of the coordinates x , consider a typical general-relativistic measurement. A standard application of GR is in precision measurements and precision modeling of solar system dynamics. In this context, partial observable quantities are the "instantaneous" distances d_p between the Earth and the different planets, defined as the proper time elapsed on Earth (measured by a clock at rest on Earth) while a radar signal goes from Earth to a planet p and back. Fixing an arbitrary initial event on Earth, one additional partial observable can be obtained as the proper time

τ from this event along the Earth trajectory. Complete observables are then the values $d_p(\tau)$ of the planet distances, at different local proper times τ . A general relativistic model of the solar system, with appropriately chosen initial data, can predict $d_p(\tau)$ for all p 's and all τ 's, and these predictions can be compared with experience. In building up this model, we choose an arbitrary coordinatization x of the solar system region, and express the gravitational field, the electromagnetic field, and the planets' positions in this coordinate system. The predicted quantities $d_p(\tau)$ are complicated non-local functions of the fields and planets' positions, which are independent of the coordinates x chosen. To be sure, the observable τ is introduced here only for convenience. We can equivalently express the predictions of the theory simply as a set of relations $f(d_p) = 0$ that must hold between the partial observables d_p .

General covariance and Kretschmann objection

The invariance of GR under active diffeomorphisms follows from two properties of the GR field equations. First, they are generally covariant. That is, they maintain the same form under any smooth change of coordinates $x \rightarrow x'(x)$ on M . This means that there is no coordinate systems on M which is preferred a priori. Second, there are no fixed non-dynamical fields in the field equations.

The first property, namely general covariance, is the property that Einstein most insisted upon, and that guided him in finding GR. The requirement of general covariance still plays a major role in selecting physical theories compatible with what we have understood about the world as a result of the general relativistic revolution.

However, general covariance *alone* is not excessively significant. Indeed, any field equation can be written in an arbitrary coordinate system. This fact was pointed out by Kretschmann shortly after Einstein wrote GR [Kretschmann, 1917], and has raised much discussion. As an example, consider the field equation

$$(11) \quad (\partial_T^2 - \partial_X^2 - \partial_Y^2 - \partial_Z^2) \varphi(X, Y, Z, T) = 0.$$

If we introduce arbitrary coordinates x (with components x^μ) as functions $x^\mu = x^\mu(X, Y, Z, T)$ (with inverse $X(x), Y(x), Z(x), T(x)$), the wave equation (11) becomes the generally covariant equation

$$(12) \quad \square_g \varphi(x) \equiv \partial_\mu \sqrt{\det -g(x)} g^{\mu\nu}(x) \partial_\nu \varphi(x) = 0,$$

In this equation, the unknown is $\varphi(x)$, while $g^{\mu\nu}(x)$ and $\det g(x)$ are the inverse and the determinant of the *fixed* field

$$(13) \quad g_{\mu\nu}(x) = \frac{\partial X(x)}{\partial x^\mu} \frac{\partial X(x)}{\partial x^\nu} + \frac{\partial Y(x)}{\partial x^\mu} \frac{\partial Y(x)}{\partial x^\nu} + \frac{\partial Z(x)}{\partial x^\mu} \frac{\partial Z(x)}{\partial x^\nu} - \frac{\partial T(x)}{\partial x^\mu} \frac{\partial T(x)}{\partial x^\nu}.$$

The field theory for the scalar field φ defined by equation (12) is *not* diffeomorphism invariant, because in distinct coordinate systems the field equations for the unknown φ are different, in the sense that they are determined by different functions $g_{\mu\nu}(x)$.

Equation (12), on the other hand, can become one of the equations of a diffeomorphism invariant theory in which $g_{\mu\nu}(x)$ is also one of the unknown, namely a dynamical field. Therefore whether or not a theory is diffeomorphism invariant is not determined just by the aspect of an equation, but by the full specification of the dynamical quantities and their equations of motion.

Einstein's insistence on general covariance alone, however, should probably not be interpreted as a lack of clarity on his part, but only as his effort to emphasize the importance of a step that was *necessary*, not sufficient, to write a successful relativistic field theory.

In addition, like all formal properties of physical theories, even full diffeomorphism invariance, should probably not be interpreted as a rigid selection principle, capable of selecting physical theories *just by itself*. With sufficient acrobatics, any theory can perhaps be re-expressed in a diffeomorphism invariant language. The same is true for any other formal invariance property. For instance, any theory can be rewritten as a rotational invariant theory, or with any other desired invariance property, by simply adding variables.

As an example, equation (11) can be viewed as physically equivalent to the diffeomorphism invariant system

$$(14) \quad \square_g \varphi(x) = 0, \quad Riem[g] = 0,$$

where $Riem[g]$ is Riemann's curvature and the unknowns are now the two fields φ and g . But there are prices to pay. First, this theory has a "fake" dynamical field, since g is constrained to a single solution up to gauges, by the second equation of the system. Having no physical degrees of freedom, g is physically a fixed background field, in spite of the trick of declaring it a variable and then constraining the variable to a single solution. Second, we can insist on a lagrangian formulation of the theory (14) [Sorkin, 2002], but to do this we must introduce an additional field, and it can then be argued that the resulting theory, having an additional field is different from (12) [Earman, 1989].

Diffeomorphism invariance is the key property of the mathematical language used to express the key conceptual shift introduced with GR: the world is not formed by a fixed non-dynamical spacetime structure, which defines localization and on which the dynamical fields live. Rather, it is formed solely by dynamical fields in interactions with one another. Localization is only defined, relationally, with respect to the fields themselves.

Relationalism and substantivalism

A non-dynamical background space was used by Newton. The first part of the *Principia*, Newton argues very explicitly that we must assume the existence of space as an entity. This part can be read as a polemic against the long dominant, and in particular Descartes's, relational understanding of space.

The two traditional views about space, absolute ("space is an entity") and relational ("space is a relation between entities"), suitably modified to take into

account scientific progress, continue in contemporary philosophy of science under the names of *substantivalism* and *relationalism*.

We can say that Einstein has “unmasked” the entity introduced by Newton (which much disturbed Leibniz): Newton’s space is nothing else than a field like the others, thought Newton considered it in a regime in which its dynamics could be neglected. Localization in space and in time, introduced by Newton against Descartes’s relational localization, is revealed by Einstein to be, after all, still a relational location — in the sense of Descartes —, with respect to a specially chosen entity: the gravitational field. In a sense, we can therefore say that GR realizes a full return to a relational definition of space and time, after the Newtonian substantivalist parenthesis.

In other words, in prerelativistic physics, spacetime is a sort of structured container which is the home of the world. In general-relativistic physics, on the other hand, there is nothing of the sort. There are only interacting fields (including the gravitational field) and particles: the only notion of localization which is present in the theory is relative: dynamical objects (fields and particles) are localized only with respect to one another. This is the notion of relational space defended by Aristotle and Descartes, against which Newton wrote the initial part of the *Principia*. Newton had two points in his favor: the physical reality of inertial effects such as the concavity of the water in the bucket of his famous bucket experiment, and the immense empirical success of his theory based on absolute space. Einstein has provided an alternative interpretation for the cause of the concavity — the interaction with an entity: the local gravitational field — and a theory based on relational space that is empirically far more effective than Newton theory. Einstein has therefore reopened the possibility of a relational understanding of space and time, which was closed by Newton’s bucket.

At the basis of Cartesian relationalism was the notion of “contiguity”. Two objects are contiguous if they are adjacent to one another. Space is the order of things with respect to such contiguity relation. At the basis of the spacetime structure of GR there is a very similar notion: Einstein’s “spacetime coincidence” is strictly analogous to Descartes’ “contiguity”.

The key to this novel relational understanding of space and time is Faraday’s revolutionary idea that a field is a physical entity. Recall that Faraday visualized a field as a family of real lines filling up everything. Einstein’s entire theoretical work fully implements this realistic interpretation of the fields. In Einstein popular-science writing, the gravitational field is a huge “jelly fish”, a better metaphor than the lines of Faraday. Entities are not just particles, but also fields, the gravitational field is one field among the others. These entities are localized only with respect to one another.

A substantivalist position can nevertheless still be –and in fact is still– defended. Einstein’s discovery that Newtonian spacetime and the gravitational field, are the same entity, can also be expressed by saying that there is no gravitational field: it is spacetime that has dynamical properties. This choice is not uncommon in the literature. The difference with the language used here is only a matter of choice

of words. The substantivalist can therefore claim use the that, according to GR, “spacetime is an entity”: indeed, it is the gravitational field, which is an entity. Since it is possible to define localization with respect to the gravitational field, the substantivalist can also say that “spacetime is an entity that defines localization”.

However, this is an extremely weakened substantivalist position. To what extent is general-relativistic spacetime different from any other arbitrary entity with respect to which we can define a relational localization? We can call “spacetime” anything used to define localization. Newton’s acute formulation of substantivalism contains a precise characterization of “space” [Newton, 1962]:

“... So it is necessary that the definition of places, and hence of local motion, be referred to some motionless thing such as extension alone or “space”, *in so far as space is seen to be truly distinct from moving bodies.*”

(My italic.) The characterizing feature of space, according to this substantivalist manifesto, is to be “truly distinct from moving” bodies. In modern terms and after the Faraday and Maxwell conceptual revolution, I believe this can only be translated as being “truly distinct from dynamical entities such as particles or fields”. This is *not* the case for the spacetime of GR. The modern substantivalist can give up Newton’s strong substantivalism (“spacetime is a non-dynamical entity”) for the much weaker thesis “we call spacetime the gravitational field, which is a dynamical entity”. But then what is the difference between this position and the relationalist one, if not just a choice of words?

To be sure, general relativistic relationalism doesn’t fit comfortably with traditional relationalism either. E.g. observables of GR, conceived as coincidence quantities are non-substantival in that they don’t require spacetime points to support them. But neither are they relational in the traditional sense of involving relations between “material” bodies or events in their histories.

The traditional substantivalist-relational alternative was formulated before the Faraday-Maxwell conceptual revolution, without taking the existence of the *fields* into account. After Faraday and Maxwell, we understand the world also in terms of a new set of dynamical entities, the fields. Once we accept the existence of the fields, and Einstein’s discovery that Newton’s space is one of the fields, the distinction between substantivalism and relationalism is largely reduced to mere semantics.

When two opposite positions in a long-standing debate have come so close that their distinction is reduced to semantics, one can perhaps say that the issue is resolved. In this sense, it may be argued that GR has solved the long-standing issue of the relational versus substantivalist interpretation of space.

4.2 Background independence

Is QM compatible with the general relativistic notions of space and time sketched above? It is, but a sufficiently general formulation of QM must be used. For instance, the Schrödinger picture is only viable for theories where there is a global

observable time variable t ; this conflicts with GR, where no such variable exists. Therefore the Schrödinger picture makes little sense in a background independent context. But formulations of QM have been proposed that are more general than the Schrödinger picture. See for instance [Hartle, 1995] and [Rovelli, 2004]. Formulations of this kind are sometimes denoted “generalized quantum mechanics”, although they might be called “quantum mechanics” in the same sense in which “classical mechanics” is used to designate formalisms with different degrees of generality, such as Newton’s, Lagrange’s, Hamilton’s or symplectic mechanics.

On the other hand, most of the conventional machinery of perturbative QFT is profoundly incompatible with the general relativistic framework. There are many reasons for this: (i) The conventional formalism of QFT relies on Poincaré invariance. In particular, it relies on the notion of energy and on the existence of the nonvanishing hamiltonian operator that generates unitary time evolution. The vacuum, for instance, is the state that minimizes the energy. But, in a general relativistic theory there is, in general, no global Poincaré invariance, no general notion of energy and no nonvanishing hamiltonian operator. (ii) At the roots of conventional QFT is the physical notion of particle. The theoretical experience with QFT on curved spacetime [Fulling, 1989] and on the relation between acceleration and temperature in QFT [Wald, 1994] indicates that in a generic gravitational situation the notion of particle can be quite delicate. (iii) Consider a conventional renormalized QFT. The physical content of the theory can be expressed in terms of its n -point functions $W(x_1, \dots, x_n)$. We expect the n -point functions to be invariant under the invariances of the theory. In a general relativistic theory, invariance under an arbitrary coordinate transformation $x \rightarrow x' = x'(x)$ implies immediately that the n -point functions must satisfy

$$(15) \quad W(x_1, \dots, x_n) = W(x'(x_1), \dots, x'(x_n)).$$

Since any set of n (distinct) points (x_1, \dots, x_n) can be transformed into any other set by a generic coordinate transformation, it follows that W is constant! It does depend on its arguments! Clearly we are in a very different framework from conventional QFT.

There is a possible escape strategy to circumvent these difficulties: write the gravitational field as the sum of two terms, as in equation (1), and assume that spacetime and causal relations are defined by the first term, rather than by the full gravitational field. This escape strategy brings back a background spacetime. A formulation of QG that does *not* take this escape strategy, and thus maintains the full symmetry of GR, is called *background independent*.

The divide

Different research directions are oriented by different evaluations given to the general relativistic spacetime conceptual revolution discussed above in Section 4.1. If this conceptual revolution is taken seriously, and understood as a feature of the world that we have learned, the problem of QG becomes the problem of understanding how to define and interpret a background independent QFT. This point

of view orients a large part of the research in loop quantum gravity and similar approaches. Not surprisingly, this line of research is more strongly influenced by the GR research tradition.

On the other hand, if the GR conceptual shift is viewed as accidental, the motivation for developing QG comes more from other open problems such as the problem of the unification (see below). One argument often presented for this point of view is that since QG affects microphysics, we can always choose a scale which is sufficiently small to disregard macroscopic curvature effects and sufficiently large to disregard QG effects. At this scale the world is Lorentz invariant. Therefore in QG we can always assume the existence of an asymptotic Lorentz invariant region. This suggests that we can use techniques associated with asymptotic Lorentz invariance. This line of thinking, predominant in the string community, is more influenced by the particle-physics tradition, which is deeply wedded to Poincaré invariance and which has mostly neglected gravity throughout the twentieth century.

The cultural divide is sometimes very strong, in spite of repeated efforts to fill the gap. Both sides feel that the other side is incapable of appreciating something basic and essential: the structure of QFT as it has been understood in half a century of investigation, for the particle-physics side; the novel physical understanding of space and time that has appeared with GR, for the relativity side. Both sides expect that the other's point of view will turn out, at the end of the day, to be not very relevant. One side because GR is only a low energy limit of a much more complex theory, and thus cannot be taken too seriously as an indication about the deep structure of Nature. The other, because the experience with QFT is on a fixed metric spacetime, and thus is irrelevant in a genuinely background independent context.

4.3 *The nature of time*

Much has been written about the fact that the main equation of nonperturbative QG, namely the Wheeler-DeWitt equation (5) does not contain the time variable t . This presentation of the “problem of time in QG”, however, is misleading, since it confuses the aspect of the problem that is specific to QG and the one which is already present in classical GR. Indeed, classical GR can be entirely formulated in the Hamilton-Jacobi formalism in terms of equation (4), where no time variable appears either.

In the classical general-relativistic context, the notion of time differs strongly from the one used in the special-relativistic context (and even more strongly from the one used in the pre-relativistic context). In the pre-relativistic context, following Newton, we assume that there is a universal physical variable t , measured by clocks, such that all physical phenomena can be described in terms of evolution equations in the independent variable t . In the special-relativistic concept, this notion of time is weakened. Clocks do not measure an universal time variable, but rather a proper time elapsed along inertial trajectories. If we fix a Lorentz frame, however, we can still describe all physical phenomena in terms of evolution

equations in the independent variable x^0 , even though this description hides the covariance of the system.

In the general relativistic context, we must distinguish two kinds of problems, that are often improperly confused. First, we can consider the problem of the dynamics of matter interacting with a given gravitational field, or, equivalently, on a given spacetime geometry. In this case, the fixed gravitational field still determines the value of the proper time τ elapsed along any (timelike) spacetime trajectory, measured by a clock moving along that trajectory. That is, a given gravitational field determines a local notion of time.

A distinct problem is given by the dynamics of the gravitational field itself, or by the interacting dynamics of gravity and matter. In this case, there is no external time variable that can play the role of observable independent evolution variable. The field equations are written in terms of an evolution parameter, which is the time coordinate x^0 , but this coordinate, as explained above in section 4.1, does not correspond to anything observable. In general, the proper time τ along spacetime trajectories also cannot be used as an independent variable, as τ is a complicated non-local function of the gravitational field itself. Therefore, properly speaking, GR does not admit a description as a system evolving in terms of an observable time variable. This is particularly evident in the Hamilton-Jacobi formulation (4) of GR. This does not mean that GR lacks predictivity. Simply put, what GR predicts are relations between partial observables, which cannot in general be represented as dependence of dependent variables on a preferred independent time variable.

To be sure, the ontological status of the time variable t is far from being straightforward in Newtonian physics either. In Newtonian physics we describe the world in terms of physical variables $A(t), B(t), \dots$ evolving in t . One may notice that in a sense we never directly access t , but only physical variables A, B, \dots , since the clock devices used to measure t are themselves physical systems with an observable time-dependent variable $C(t)$, such as the position of the clock's hand. Therefore, what we actually observe is always the relative evolution of observable variables $A(C), B(C), A(B) \dots$ and never t itself. Newton makes this point clearly in the *Principia*, but also observes that the direct mathematicization of the apparent motions $A(C), B(C), A(B) \dots$ becomes greatly simplified by hypostatizing the existence of t , and expressing all evolution in terms of t . This of course works excellently in the nonrelativistic and nongravitational context. But it is not illogical that Newton's strategy might fail in certain regimes. And in fact it fails in the relativistic gravitational regime, where no universal t can be introduced, and we can only describe the relative dependence of observable quantities. This is what happens in GR.

In a sense, any partial observable variable can be chosen as the independent one in GR. In general, none has the idealized properties assumed by the Newtonian time t , which grows monotonically irrespectively of the state of the system. For instance, in a closed cosmology the volume a of the universe and the proper time t_c since the Big Bang, along a galaxy worldline, are often used as independent

variables. But a behaves badly if the Universe begins recontracting, and t_c is only defined in the approximation in which the Universe is assumed to be homogeneous (what is the value of t_c if two galaxies with different proper time from the big bang meet?)

Such a weakening of the notion of time in classical GR is rarely emphasized, because, after all, in classical physics we may disregard the full dynamical structure of the dynamical theory and consider only a single solution of its equations of motion. As mentioned, a single solution of the GR equations of motion determines a spacetime, where a notion of proper time is associated to each timelike worldline. In the quantum context, on the other hand, there is no single spacetime, as there is no trajectory for a quantum particle, and the very concept of time becomes fuzzy.

Attitudes towards the problem of time

Different attitudes can be found in the literature with regard to the problem of time. For technical overviews and references (not completely up to date), see for instance [Isham, 1992; Kuchar, 1992].

As already mentioned, a considerable part of QG research disregards the issue, and maintains that Minkowski space, Poincaré invariance, with its associated notion of time evolution (as a subgroup of the Poincaré group), should not be abandoned in building QG, notwithstanding the features of GR.

Other authors maintain that even if Poincaré invariance is lost in GR and the notion of time becomes more complex, still the idea that the world exists in time, and that its description is the description of systems evolving in time, is a primary notion that we cannot renounce.

Some of these authors have proposed minor modifications of GR, capable of reintroducing a fundamental notion of observable time evolution in the theory. One possibility is to choose a preferred gauge-fixing, in which diffeomorphism invariance is partially broken, and the time coordinate is gauge fixed to be equal to some function of the gravitational field. An example is York time, defined as the trace of the extrinsic curvature of a spacelike surface. Alternatively, the dynamics of GR can be modified, to get a theory with an independent time parameter.

Others accept in full the challenge presented by GR of trying to conceptualize the world in the absence of a fundamental notion of time and time evolution, as illustrated in the following section.

Physics without space and time?

An illustrative example of how a formulation of mechanics might not use space and time as independent variables is provided by the following proposal (see [Rovelli, 2004], Chapters 4 and 6). Consider a finite spacetime region R bounded by a closed three-dimensional surface Σ . Let (φ, g) represent the value of all fields, including the gravitational field g , on Σ , and let (P_φ, P_g) represent the normal derivative of the fields out of Σ . In principle, all predictions of classical GR can be expressed as

constraints on the possible values that the set $(\varphi, g), (P_\varphi, P_g)$ can take. Similarly, in principle all the predictions of QG can be expressed in terms of the probability amplitude $W(\varphi, g)$ of measuring the fields (φ, g) . This is a generalization of Feynman's observation that the quantum dynamics of a particle is contained in the propagator $W(x', t'; x, t)$. Diffeomorphism invariance implies that $W(\varphi, g)$ does not depend on the way Σ is imbedded into M . In other words, the entire quantitative spatial and temporal dependence is encoded into the dependence of $W(\varphi, g)$ on the gravitational field g . If, for instance, we identify Σ with the surface of the initial, final and boundary values of a scattering experiment, then it is only the value of the gravitational field on the boundary that determines the time lapses between the initial and final surfaces. Recall indeed that in GR spatial distances and temporal intervals are functions of the gravitational field. $W(\varphi, g)$ can then be used in principle to determine all possible probabilistic predictions regarding the experiment, without using independent spatial or temporal variables.

In order to understand the quantum gravitational field, some of the emphasis on geometry should probably be abandoned. Geometry represents well the classical gravitational field, not quantum spacetime. This is not a betrayal of Einstein's legacy: on the contrary, it is a step in the direction of "relativity" in the precise sense meant by Einstein. The key conceptual difficulty of QG may therefore to find a way to understand the physical world in the absence of the familiar stage of space and time. What might be needed is to free ourselves from the prejudices associated with the habit of thinking of the world as "inhabiting space" and "evolving in time".²

Whether it is logically possible to understand the world in the absence of fundamental notions of time and time evolution, and whether this may be consistent with our experience of the world is an open question.

Unitarity

Absence of a fundamental notion of time evolution implies in particular that there is no unitary time evolution in the theory. Absence of unitarity is viewed with great suspicion by many physicists coming from the high energy tradition, where the requirement of unitarity has repeatedly played a major historical role. The argument is often put forward that a probabilistic theory without unitary time evolution is inconsistent. This is not correct, since inconsistency follows from lack of unitarity in the presence of a standard time evolution, and not in the absence of it. If, for instance, we describe the evolution of the universe using the volume of

²If we take this extreme attitude, one problem is to recover the macroscopic notion of time evolution and the specific features of the macroscopic time observable, from an atemporal microscopic theory. It is well known that it is surprisingly hard to pin-point with precision what characterizes the time variable in a dynamical system; on the other hand, the thermodynamical and statistical behavior of physical systems is strongly temporally characterized. Accordingly, the hypothesis has been considered [Rovelli, 1993a; Rovelli, 1993b; Connes and Rovelli, 1994] that "temporal flow" is a feature of the world that appears only in the context of a statistical-thermodynamical description. In other words, "time" could be an artifact of our vast ignorance of the microstate of the world.

the universe a as independent variable, there is no reason to require the probability for the universe to exist to be unit at all a . Indeed, there is a finite probability that the universe reach only a maximum value of a and then re contracts. On the other hand, the consistency of a probabilistic interpretation of QM in a context in which evolution is not expressed in terms of an external variable t is still unclear.

5 RELATION WITH OTHER OPEN PROBLEMS

In the history of physics, often two open problems have found a common solution. For instance, the problem of understanding the nature of light and the problem of unifying the electric and magnetic theory found a common solution in Maxwell theory. Often, however, the hope to solve two problems at once has been disappointed. For instance, in the sixties the hope was strong to find a theory for the strong interaction and at the same time get rid of renormalization theory; but QCD turned out to be a good solution of the first problem without addressing the second. The problem of QG has been suggested to be related to all sorts of open problems in theoretical physics.

5.1 *Unification*

The current description of the physical world is composed by a number of field theories: on the one hand GR, on the other hand the standard model which in turn is composed of the electroweak theory and quantum chromodynamics; in addition, fermions are present in several multiplets, and there are the Higgs scalars. The theory has more than a dozen elementary constants. In the wake of the successful unifications of electric and magnetic theory, and then of the electromagnetic and the weak interactions theories, research has long aimed to reduce the complexity of the standard model by providing a single coherent theory governed by a smaller number of elementary constants.

Opinions diverge on the relation between this “unification” problem and the problem of QG. *A priori*, there is no strict reason why the quantum properties of gravity should be understood only in conjunction with the other field theories; the quantum properties of electromagnetism, for instance, have been understood in the context of QED without reference to the other interactions, and so have the properties of the strong interactions.

Some arguments have been proposed to support the idea that the two problems must be solved together. I mention three: the first is speculative and I think weak. The second and third are technical and have some weight. First, there is a widespread expectation that a final “Theory Of Everything” should be at hand today. Historically, however, this expectation has been often present in theoretical physics, and so far always erroneously.

The second argument comes from the early history of the attempts at replacing GR with a renormalizable theory: supergravity has shown that the gravitational ultraviolet divergences are suppressed (although, at the end of the day, not cured)

by an appropriate coupling between gravity and matter (a fermion field, in the case of supergravity).

The third argument supporting the relation between the two problems is the following. In the standard model, the coupling constants that determine the strength of the electromagnetic, weak and strong interactions depend on the scale of the phenomena considered. At normal scales, they are widely different in size, but they converge at a scale which is quite close to the Planck scale. This suggests that the scale at which unification might take place should be the same as the scale at which quantum gravitational effects become manifest, indicating that the two phenomena are likely to be related.

Concretely, QG is realized in string theory in the context of a tight unification, while loop quantum gravity proposes a solution of the QG problem unrelated to unification.

5.2 *Interpretation of quantum mechanics*

In spite of its enormous empirical success and its nowadays ubiquitous applications, QM is a theory which is viewed by many as not yet completely understood. The interpretation of the theory is relatively uncontroversial as long as we use it to describe physical systems interacting with an external system (the “observer”) whose quantum properties can be disregarded. But a number of difficulties appear as soon as we take the quantum properties of the observer into account. In the physics community, the attitudes to this problem vary widely, ranging from a complete denial that a problem exists to various proposals for modifying QM in order to solve it. But the number of physicists who consider this a genuine open problem has been increasing in the last decade.

Various arguments have been proposed to tie the problem of the interpretation of QM to the QG problem. One is, once more, the expectation that a final theory might be at hand, and the final theory must be entirely self consistent.

Roger Penrose has proposed a specific mechanism via which quantum linearity might be broken by gravity: gravity might be a physical factor inducing a physically realized wave-function collapse [Penrose, 1986]. The proposal is in principle empirically testable.

In the context of Smolin’s and Adler’s attempts to derive quantum mechanics from the statistical behavior of a statistical dynamics of matrix models [Smolin, 2002; Adler, 2004], the suggestion has been made that one might seek a common origin for both gravitation and quantum field theory at a common deeper level of physical phenomena from which quantum field theory emerges.

Finally, the suggestion has been made that the relational aspect of spatiotemporal structure revealed by GR could be connected with the relational aspect of QM emphasized by the “relational” interpretations of QM [Rovelli, 1996b]. The first is determined by the relation of contiguity between systems; the second by the interaction between systems. But on the one hand locality implies that interaction happens only between contiguous systems, and on the other hand contiguity

is only manifest via a physical interaction, suggesting a strict connection between the two relations. These ideas however, have not been developed beyond the stage of suggestions.

5.3 *The cosmological constant*

An elusive aspect of the current description of the universe is given by the cosmological constant, a constant introduced by Einstein, describing a long range gravitational coupling that can modify gravity at large distances. This constant plays a major role in cosmology, in QFT (where quantum field theoretical effects tend to make it unrealistically large), and recent cosmological observations seem to indicate that its value is very small but not, as previously expected, vanishing.

Although nothing clear has so far appeared in QG research concerning this constant, it must be noted that Raphael Sorkin's QG theory predicted a small value of the constant with the correct order of magnitude, before its observation, as mentioned in Section 2.1.7.

5.4 *Quantum cosmology*

“Quantum cosmology” indicates the study of the Universe as a whole as a quantum system. There are two distinct problems that go under this name.

The first is the quantum version of the modelling of the dynamics of our Universe: in particular, the study of the quantum features of the dynamical systems obtained under the drastic simplification that the Universe is homogenous. These classical models, such as the Friedmann-Robertson-Walker model, play an important role in cosmology and are believed to give a good description of the large scale features of our Universe. Their quantization is of interest on several grounds. First, it provides a simplified framework in which many of the conceptual difficulties of QG can be examined and solutions can be tested. Second, they can be used to study what a quantum theory of gravity could us concerning the physics near and at the Big Bang itself, where quantum gravitational effects are expected to dominate.

The study of these models has been started in the sixties by Bryce DeWitt [DeWitt, 1967b; DeWitt, 1967c] and Charles Misner [Misner, 1969] and has seen a great development in the following decades. The limitation of these models, of course, is that they are based on the freezing of all the infinite numbers of degrees of freedom of GR, except for a finite number of them, and therefore they miss the entire field theoretical aspect of the QG problem.

A string cosmology has been developed by Gabriele Veneziano and collaborators, with the hope of finding observational consequences of string theory [Gasperini and Veneziano, 1993]. The application of loop quantum gravity to quantum cosmology (“loop cosmology”) has recently led to a model which is finite and well-behaved at the initial singularity [Bojowald, 2001; Bojowald and Morales-Tecotl, 2006].

The second problem that goes under the name of “quantum cosmology” is the conceptual problem of describing a quantum system that forms the entire universe, and therefore for which there is no “external” observer: i.e., the study of quantum mechanics in the case in which the observer is inside the system.

This second problem is very loosely related to the problem of quantum gravity. It is true that it is impossible to be “external” with respect to the gravitational field, but one should not confuse “external” in the spatiotemporal sense with “external” in the dynamical sense. One cannot be “external” with respect to the electromagnetic field either, in the spatiotemporal sense; but we can nevertheless consider an electromagnetic system, viewed as a quantum system, interacting with an external system, viewed as a classical observer. The same can be done for a gravitational system. Therefore nothing *a priori* prevents us from using the standard Copenhagen interpretation of QM (whether or not this is satisfactory) in the context of QG. In other words, the problem of QG and this second problem of quantum cosmology are not necessarily related.

On the other hand, the difficulties raised by considering the observer as part of the system and the difficulties generated in QM by diffeomorphism invariance, in particular the absence of an external time, are of a similar nature, and both question the viability of the Copenhagen interpretation. A general scheme for addressing both kinds of difficulties, and defining a generalized formalism for QM, where there is no external time and no external observer, has been developed by Jim Hartle [Hartle, 1995].

6 CONCLUSION

After 70 years of research, there is no consensus, no established theory, and no QG theory has yet received any direct or indirect experimental support. In the course of 70 years, many ideas have been explored, fashions have come and gone, the discovery of the Holy Grail of QG has been several times announced, only to be later greeted by much scorn.

However, in spite of this, research in QG has not been meandering meaninglessly. On the contrary, a consistent logic has guided the development of the research, from the early formulation of the problem and the research directions in the fifties to nowadays. The implementation of the programs has been laborious, but has been achieved. Difficulties have appeared, and solutions have been proposed, which, after much difficulty, have led to the realization, at least partial, of the initial hopes.

It was suggested in the early seventies that GR could perhaps be seen as the low energy limit of a theory without uncontrollable divergences; today, 30 years later, such a theory — string theory — is known. In 1957 Charles Misner indicated that in the canonical framework one should be able to compute eigenvalues; and in 1995, 37 years later, eigenvalues were computed — within loop quantum gravity. Much remains to be understood and some of the current developments might lead nowhere. We are not at the end of the road, we are only half-way through the

woods. But looking at the entire development of the subject, it is difficult to deny that there has been substantial progress.

The progress cannot be just technical. The search for a quantum theory of gravity raises again old questions such as: What is space? What is time? What is the meaning of “being somewhere”? What is the meaning of “moving”? Is motion to be defined with respect to objects or with respect to space? Can we formulate physics without referring to time or to spacetime? And also: What is causality? What is the role of the observer in physics?

Questions of this kind have played a central role in periods of major advances in physics. For instance, they played a central role for Einstein, Heisenberg, Bohr and their colleagues. But also for Descartes, Galileo, Newton and their contemporaries, and for Faraday, Maxwell and their colleagues. Today some physicists view this manner of posing problems as “too philosophical”. Most physicists of the second half of the twentieth century, indeed, have viewed questions of this nature as irrelevant. This view was appropriate for the problems they were facing. When the basics are clear and the issue is problem-solving within a given conceptual scheme, there is no reason to worry about foundations: a pragmatic approach is the most effective one. Today the kind of difficulties that fundamental physics faces has changed. To understand quantum spacetime, physics has to return, once more, to those foundational issues. We have to find new answers to the old foundational questions. The new answers have to take into account what we have learned with QM and GR. The problem of QG will probably not be solved unless these questions are carefully reconsidered.

ACKNOWLEDGEMENTS

Thanks to Christian Wüthrich, Jeremy Butterfield and John Earman for their careful reading of the manuscript and very valuable suggestions.

BIBLIOGRAPHICAL NOTE

For more details on the history of QG see the historical appendix in [Rovelli, 2004]; and, for early history see [Stachel, 1999b; Stachel, 1999a] and [Gorelik, 1992]. For orientation on current research on QG, see the review papers [Horowitz, 2000; Carlip, 2001; Isham, 1991; Rovelli, 1998a]. An interesting panorama of points of view on the problem and on philosophical issues it raises is in the various contributions to the book [Callender and Huggett, 2001]. See also the discussion in [Rovelli, 1997; Rovelli, 2000]. As a general introduction to QG ideas, see the old classic reviews, which are rich in ideas and present different points of view, such as John Wheeler 1967 [Wheeler, 1968], Steven Weinberg 1979 [Weinberg, 1979], Stephen Hawking 1979 and 1980 [Hawking, 1979; Hawking, 1984], Karel Kuchar 1980 [Kuchar, 1984], and Chris Isham’s magisterial syntheses [Isham, 1984a; Isham, 1984b; Isham, 1997]. On string theory, classic textbooks are Green, Schwarz and Witten, and

Polchinski [Green *et al.*, 1987; Polchinski, 1998]. On loop QG, see [Rovelli, 1998b; Rovelli, 2004]. For a discussion of the difficulties of string theory and a comparison of the results of strings and loops, see [Rovelli, 2003], written in the form of a dialogue, and [Smolin, 2003]. Smolin's popular book [Smolin, 2000] provides a readable introduction to QG.

BIBLIOGRAPHY

- [Adler, 2004] S. Adler. *Quantum Theory as an Emergent Phenomenon*. Cambridge University Press, 2004.
- [Amati *et al.*, 1989] D. Amati, M. Ciafaloni, and G. Veneziano. Can spacetime be probed below the string size? *Physics Letters B*, 216:41–47, 1989.
- [Ambjorn *et al.*, 1997] J. Ambjorn, M. Carfora, and A. Marzuoli. *The Geometry of Dynamical Triangulations*. Lecture Notes in Physics. Springer-Verlag, Berlin, 1997.
- [Arnowitt *et al.*, 1962] R. Arnowitt, S. Deser, C. Misner. The dynamics of general relativity. In L. Witten (ed.), *Gravitation: An Introduction to Current Research*, page 227. Wiley, New York, 1962.
- [Ashtekar, 1986] A. Ashtekar. New variables for classical and quantum gravity. *Physical Review Letters*, 57:2244–2247, 1986.
- [Ashtekar, 1987] A. Ashtekar. New Hamiltonian formulation of general relativity. *Physical Review D*, 36:1587–1602, 1987.
- [Ashtekar *et al.*, 1998] A. Ashtekar, J. Baez, A. Corichi, and K. Krasnov. Quantum geometry and black hole entropy. *Physical Review Letters*, 80:904–907, 1998. gr-qc/9710007.
- [Bardeen *et al.*, 1973] J. Bardeen, B. Carter, and S. Hawking. The four laws of black hole mechanics. *Communication in Mathematical Physics*, 31:161–170, 1973.
- [Bekenstein, 1972] J. Bekenstein. Black holes and the second law. *Nuovo Cimento Letters*, 4:737–740, 1972.
- [Bekenstein, 1973] J. Bekenstein. Black holes and entropy. *Physical Review*, D7:2333–2346, 1973.
- [Bekenstein, 1974] J. Bekenstein. Generalized second law for thermodynamics in black hole physics, 1974. *Physical Review*, D9:3292–3300.
- [Belot, 1998] G. Belot. Why general relativity does need an interpretation. *Philosophy of Science*, 63:S80–S88, 1998.
- [Bergmann, 1949a] P. Bergmann. Non-linear field theories. *Physical Review*, 75:680–685, 1949a.
- [Bergmann, 1949b] P. Bergmann. Non-linear field theories II: Canonical equations and quantization. *Review of Modern Physics*, 21:480–487, 1949b.
- [Bergmann, 1958] P. Bergmann. Conservation laws in general relativity as the generators of coordinate transformations. *Physical Review*, 112:287–289, 1958.
- [Bergmann, 1961] P. Bergmann. Observables in general relativity. *Review of Modern Physics*, 33:510–514, 1961.
- [Bergmann and Komar, 1980] P. Bergmann and A. Komar. The phase space formulation of general relativity and approaches towards quantization. In A. Held (ed.), *General Relativity and Gravitation: One hundred years after the birth of Einstein*, pages 227–254. Plenum Press, New York, 1980.
- [Bojowald, 2001] M. Bojowald. Absence of singularity in loop quantum cosmology. *Physical Review Letters*, 86:5227–5230, 2001. gr-qc/0102069.
- [Bojowald and Morales-Tecotl, 2006] M. Bojowald, H. Morales-Tecotl. Cosmological applications of loop quantum gravity. gr-qc/0306008. to appear in the proceedings of the Fifth Mexican School (DGFM): *The Early Universe and Observational Cosmology.*, 2006.
- [Bronstein, 1936] M. Bronstein. Quantentheories schwacher Gravitationsfelder. *Physikalische Zeitschrift der Sowietunion*, 9:140, 1936.
- [Callender and Huggett, 2001] C. Callender and N. Huggett (eds.). *Physics Meets Philosophy at the Planck scale*. Cambridge University Press, Cambridge, 2001.
- [Carlip, 2001] S. Carlip. Quantum gravity: A progress report. *Reports on Progress in Physics*, 64:885, 2001. gr-qc/0108040.

- [Chamseddine and Connes, 1996] A. Chamseddine and A. Connes. Universal formula for non-commutative geometry actions: Unification of gravity and the standard model. *Physical Review Letters*, 24:4868–4871, 1996.
- [Chamseddine and Connes, 1997] A. Chamseddine and A. Connes. The spectral action principle. *Communication in Mathematical Physics*, 186:731–750, 1997.
- [Connes, 1994] A. Connes. *Non Commutative Geometry*. Academic Press, New York, 1994.
- [Connes and Rovelli, 1994] A. Connes and C. Rovelli. Von Neumann algebra automorphisms and time versus thermodynamics relation in general covariant quantum theories. *Classical and Quantum Gravity*, 11:2899–2917, 1994.
- [Deser and van Nieuwenhuizen, 1974a] S. Deser and P. van Nieuwenhuizen. One loop divergences of the quantized Einstein-Maxwell fields. *Physical Review D*, 10:401–410, 1974a.
- [Deser and van Nieuwenhuizen, 1974b] S. Deser and P. van Nieuwenhuizen. Non-renormalizability of the quantized Dirac-Einstein system. *Physical Review*, D10:411–420, 1974b.
- [DeWitt, 1964a] B. DeWitt. In Gauthier-Villars (ed.), *Conférence internationale sur les théories relativistes de la gravitation*. Editions Scientifiques de Pologne, Warsaw, 1964a.
- [DeWitt, 1964b] B. DeWitt. Theory of radiative corrections for non-Abelian gauge fields. *Physical Review Letters*, 12:742–746, 1964b.
- [DeWitt, 1965] B. DeWitt. *Dynamical Theory of Groups and Fields*. Wiley, New York, 1965.
- [DeWitt, 1967a] B. DeWitt. Quantum theory of gravity. I. The canonical theory. *Physical Review*, 160:1113–1148, 1967a.
- [DeWitt, 1967b] B. DeWitt. Quantum theory of gravity. II. The manifestly covariant theory. *Physical Review*, 162:1195–1239, 1967b.
- [DeWitt, 1967c] B. DeWitt. Quantum theory of gravity. III. Applications of the covariant theory. *Physical Review*, 162:1239–1256, 1967c.
- [Dirac, 1950] P. Dirac. Generalized Hamiltonian dynamics. *Canadian Journal of Physics*, 2:129–148, 1950.
- [Dirac, 1958] P. Dirac. The theory of gravitation in Hamiltonian form. *Proceedings of the Royal Society of London*, A246:333–343, 1958.
- [Dirac, 1959] P. Dirac. Fixation of coordinates in the Hamiltonian theory of gravitation. *Physical Review*, 114:924–930, 1959.
- [Dirac, 1964] P. Dirac. *Lectures on Quantum Mechanics*. Belfer Graduate School of Science, Yeshiva University, New York, 1964.
- [Doplicher, 1996] S. Doplicher. Quantum spacetime. New problems in the general theory of fields and particles, part II. *Annales de l'Institut Henri Poincaré*, 64:543–553, 1996.
- [Doplicher et al., 1994] S. Doplicher, K. Fredenhagen, J. Roberts. Spacetime quantization induced by classical gravity. *Physics Letters B*, 331:39–44, 1994.
- [Doplicher et al., 1995] S. Doplicher, K. Fredenhagen, J. Roberts. The quantum structure of spacetime at the Planck scale and quantum fields. *Communications in Mathematical Physics*, 172:187–220, 1995.
- [Earman,] J. Earman. Private communication.
- [Earman, 1987] J. Earman, J. Norton. What price spacetime substantivalism? The hole story. *British Journal for the Philosophy of Science*, 38:515–525, 1987.
- [Earman, 1989] J. Earman. *World Enough and Space-time: Absolute versus Relational Theories of Spacetime*. MIT Press, Cambridge, 1989.
- [Earman, 2001] J. Earman and G. Belot. Pre-Socratic quantum gravity. In C. Callender (ed.), *Physics meets philosophy at the Planck scale*. Cambridge University Press, Cambridge, 2001.
- [Einstein, 1916] A. Einstein. Naheungsweise integration der Feldgleichungen der gravitation. *Preussische Akademie der Wissenschaften (Berlin) Sitzungsberichte*, pages 688–696, 1916.
- [Einstein and Grossmann, 1914] A. Einstein and M. Grossmann. Entwurf einer verallgemeinerten Relativitätstheorie und einer Theorie der Gravitation. *Zeitschrift für Mathematik und Physik*, 62:225–261, 1914.
- [Faddeev and Popov, 1967] L. Faddeev and V. Popov. Feynman diagrams for the Yang-Mills field. *Physics Letter B*, 25:29–30, 1967.
- [Feynman, 1963] R. Feynman. Quantum theory of gravitation. *Acta Physical Polonica*, 24:697–722, 1963.
- [Fierz, 1939] M. Fierz. Ber die relativistische Theorie kräftefreier Teilchen mit beliebigem Spin. *Helvetica Physica Acta*, 12:3–37, 1939.
- [Finkelstein, 1997] D. Finkelstein. *Quantum Relativity*. Springer, Berlin, 1997.

- [Freedman *et al.*, 1976] D. Freedman, S. Ferrara, and P. van Nieuwenhuizen. Progress toward a theory of supergravity. *Physical Review*, D13:3214–3218, 1976.
- [Fulling, 1989] S. Fulling. *Aspects of Quantum Field Theory in Curved Spacetime*. Cambridge University Press, Cambridge, 1989.
- [Gasperini and Veneziano, 1993] M. Gasperini and G. Veneziano. Pre-bigbang in string cosmology. *Astroparticle Physics*, 1:317–339, 1993.
- [Gorelik, 1992] G. Gorelik. First steps of quantum gravity and the Planck values. In J. Eisenstaedt and A. Kox (eds.), *Studies in the history of general relativity*, volume 3 of *Einstein Studies*, pages 364–379. Birkhaeuser, Boston, 1992.
- [Gorelik and Frenkel, 1994] G. Gorelik and V. Frenkel. *Matvei Petrovic Bronstein and the Soviet Theoretical Physics in the Thirties*. Birkhauser Verlag, Boston, 1994.
- [Goroff and Sagnotti, 1985] M. H. Goroff, A. Sagnotti. Quantum gravity at two loops. *Physics Letter B*, 160:81–86, 1985.
- [Goroff and Sagnotti, 1986] M. H. Goroff, A. Sagnotti. The ultraviolet behaviour of Einstein gravity. *Nuclear Physics B*, 266:709–736, 1986.
- [Green and Schwarz, 1984] M. Green and J. Schwarz. Anomaly cancellation in supersymmetric $d=10$ gauge theory requires $SO(32)$. *Physics Letters B*, 149:117–122, 1984.
- [Green *et al.*, 1987] M. Green, J. Schwarz, and E. Witten. *Superstring Theory*. Cambridge University Press, Cambridge, 1987.
- [Gupta, 1952] S. Gupta. Quantization of Einstein’s gravitational field: General treatment. *Proceedings Physical Society A*, 65:608–619, 1952.
- [Hall and Hall, 1962] A. Hall and M. Hall (eds.). *Unpublished papers of Isaac Newton*. Cambridge University Press, Cambridge, 1962.
- [Hartle, 1995] J. Hartle. Spacetime quantum mechanics and the quantum mechanics of spacetime. In B. Julia, J. Z.-J. (eds.), *Proceedings on the 1992 Les Houches School, Gravitation and Quantisation*, page 285. Elsevier Science, Paris, 1995.
- [Hartle and Hawking, 1983] J. Hartle and S. Hawking. Wave function of the universe. *Physical Review*, D28:2960–2975, 1983.
- [Hawking, 1974] S. Hawking. Black hole explosions? *Nature*, 248:30–31, 1974.
- [Hawking, 1975] S. Hawking. Particle creation by black holes. *Communication in Mathematical Physics*, 43:199–220, 1975.
- [Hawking, 1979] S. Hawking. The path-integral approach to quantum gravity. In S. Hawking and W. Israel (eds.), *General Relativity: An Einstein Centenary Survey*. Cambridge University Press, Cambridge, 1979.
- [Hawking, 1984] S. Hawking. Quantum cosmology. In B. DeWitt and R. Stora (eds.), *Relativity, Groups and Topology, Les Houches Session XL*. North Holland, Amsterdam, 1984.
- [Heisenberg, 1938] W. Heisenberg. Die Grenzen der Anwendbarkeit der bisherigen Quantentheorie. *Zeitschrift für Physik*, 110:251–266, 1938.
- [Horowitz, 2000] G. Horowitz. Quantum gravity at the turn of the millenium. Plenary talk at the Marcell Grossmann conference, Rome, 2000. gr-qc/0011089.
- [Howard and Stachel, 1989] D. Howard and J. Stachel (eds.). *Einstein and the History of General Relativity*, volume I of *Einstein Studies*. Birkhauser, Boston, 1989.
- [Isham, 1984a] C. Isham. In B. DeWitt and R. Stora (eds.), *Relativity Groups and Topology. Les Houches 1983*. North Holland, Amsterdam, 1984a.
- [Isham, 1984b] C. Isham. Quantum gravity: an overview. In *Oxford 1980, Proceedings, Quantum Gravity 2*. Oxford University Press, Oxford, 1984b.
- [Isham, 1992] C. Isham. Canonical quantum gravity and the problem of time. Lectures presented at the NATO Advanced Institute: “Recent problems in Mathematical Physics, Salamanca June 15”, 1992.
- [Isham, 1994] C. Isham. Quantum logic and the histories approach to quantum theory. *Journal of Mathematical Physics*, 35:2157–2185, 1994. gr-qc/9308006.
- [Isham, 1997] C. Isham. Structural problems facing quantum gravity theory. In M. Francaviglia, G. Longhi, L. Lusanna, and E. Sorace (eds.), *Proceedings of the 14th International Conference on General Relativity and Gravitation*, pages 167–209. World Scientific, Singapore, 1997.
- [Isham, 1991] C. J. Isham. Conceptual and geometrical problems in quantum gravity. In H. Mitter and H. Gausterer (eds.), *Recent Aspects of Quantum Fields: Proceedings of the Xxx Int. Universitätswochen Fur Kernphysik, Schladming, Austria February and March 1991 (Lectur)*, page 123. Springer-Verlag, Berlin, 1991.

- [Jacobson, 1995] T. Jacobson. Thermodynamics of spacetime: The Einstein equation of state. *Physical Review Letters*, 75:1260–1263, 1995.
- [Jacobson and Smolin, 1988] T. Jacobson and L. Smolin. Nonperturbative quantum geometries. *Nuclear Physics B*, 299:295–345, 1988.
- [Klein, 1927] O. Klein. Zur fünfdimensionalen Darstellung der Relativitätstheorie. *Zeitschrift für Physik*, 46:188–208, 1927.
- [Krasnov, 1997] K. Krasnov. Geometrical entropy from loop quantum gravity. *Physical Review D*, 55:3505–3513, 1997.
- [Kretschmann, 1917] E. Kretschmann. Über den physikalischen Sinn der Relativitätspostulate. *Annals of Physics (Leipzig)*, 53:575–614, 1917.
- [Kuchar, 1984] K. Kuchar. Canonical methods of quantization. In *Oxford 1980, Proceedings, Quantum Gravity 2*. Oxford University Press, Oxford, 1984.
- [Kuchar, 1992] K. Kuchar. Time and interpretations of quantum gravity. In G. Kunstatter, D. Vincent, and J. Williams (eds.), *Proceedings of the 4th Canadian conference on General Relativity and Relativistic Astrophysics*, Singapore. World Scientific, 1992.
- [Misner, 1957] C. Misner. Feynman quantization of general relativity. *Review of Modern Physics*, 29:497–509, 1957.
- [Misner, 1969] C. Misner. Quantum cosmology. I. *Physical Review*, 186:1319–1327, 1969.
- [Newton, 1962] I. Newton. De gravitatione et aequipondio fluidorum. Translation in [Hall and Hall, 1962], 1962.
- [Norton, 1984] J. Norton. How Einstein found his field equations: 1912–1915. *Historical Studies in the Physical Sciences*, 14:253–316, 1984. Reprinted in [Howard and Stachel, 1989, 101–159].
- [Pauli and Fierz, 1939] W. Pauli and M. Fierz. On relativistic field equations of particles with arbitrary spin in an electromagnetic field. *Helvetica Physica Acta*, 12:297–300, 1939.
- [Pauri and Vallisneri, 2002] M. Pauri and M. Vallisneri. Ephemeral point-events: is there a last remnant of physical objectivity? *Dialogos*, 79:263–303, 2002.
- [Penrose, 1967] R. Penrose. Twistor algebra. *Journal of Mathematical Physics*, 8:345–366, 1967.
- [Penrose, 1986] R. Penrose. Gravity and state vector reduction. In R. Penrose and C. Isham (eds.), *Quantum Concepts in Space and Time*, page 129. Clarendon Press, Oxford, 1986.
- [Peres, 1962] A. Peres. On Cauchy’s problem in general relativity. *Nuovo Cimento*, 26:53–61, 1962.
- [Perez, 2003] A. Perez. Spin foam models for quantum gravity. *Classical and Quantum Gravity*, 20:R43–R104, 2003. gr-qc/0301113.
- [Polchinski, 1995] J. Polchinski. Dirichlet branes and Ramon-Ramon charges. *Physical Review Letters*, 75:4724–4727, 1995.
- [Polchinski, 1998] J. Polchinski. *String Theory*. Cambridge University Press, Cambridge, 1998.
- [Rosenfeld, 1930a] L. Rosenfeld. Über die Gravitationswirkungen des Lichtes. *Zeitschrift für Physik*, 65:589–599, 1930a.
- [Rosenfeld, 1930b] L. Rosenfeld. Zur quantelung der Wellenfelder. *Annalen der Physik*, 5:113–152, 1930b.
- [Rovelli, 1993a] C. Rovelli. Statistical mechanics of gravity and thermodynamical origin of time. *Classical and Quantum Gravity*, 10:1549–1566, 1993a.
- [Rovelli, 1993b] C. Rovelli. The statistical state of the universe. *Classical and Quantum Gravity*, 10:1567–1578, 1993b.
- [Rovelli, 1996a] C. Rovelli. Black hole entropy from loop quantum gravity. *Physical Review Letters*, 14:3288–3291, 1996a.
- [Rovelli, 1996b] C. Rovelli. Relational quantum mechanics. *International Journal of Theoretical Physics*, 35:1637–1678, 1996b.
- [Rovelli, 1997] C. Rovelli. Halfway through the woods. In J. Earman and J. Norton (eds.), *The Cosmos of Science*. University of Pittsburgh Press and Universitäts Verlag-Konstanz, 1997.
- [Rovelli, 1998a] C. Rovelli. Strings, loops and the others: a critical survey on the present approaches to quantum gravity. In N. Dadhich and J. Narlikar (eds.), *Gravitation and Relativity: At the turn of the Millenium (15th International Conference on General Relativity and Gravitation)*, pages 281–331. Inter-University centre for Astronomy and Astrophysics, Pune, 1998a. gr-qc/9803024.
- [Rovelli, 1998b] C. Rovelli. Loop Quantum Gravity. *Living Reviews in Relativity*, 1998b, www.livingreviews.org/Articles/Volumel/1998-Irovelli.

- [Rovelli, 2000] C. Rovelli. The century of the incomplete revolution: searching for general relativistic quantum field theory. *Journal of Mathematical Physics*, 41:3776–3800, 2000. hep-th/9910131.
- [Rovelli, 2003] C. Rovelli. A dialog on quantum gravity. *International Journal of Modern Physics D*, 12:1509–1528, 2003. hep-th/0310077.
- [Rovelli, 2004] C. Rovelli. *Quantum Gravity*. Cambridge University Press, Cambridge, 2004.
- [Rovelli and Smolin, 1988] C. Rovelli and L. Smolin. Knot theory and quantum gravity. *Physical Review Letters*, 61:1155, 1988.
- [Rovelli and Smolin, 1990] C. Rovelli and L. Smolin. Loop space representation for quantum general relativity. *Nuclear Physics B*, 331:80, 1990.
- [Rovelli and Smolin, 1995] C. Rovelli and L. Smolin. Discreteness of Area and Volume in Quantum Gravity. *Nuclear Physics B*, 442:593–622, 1995.
- [Rovelli and Speziale, 2003] C. Rovelli and S. Speziale. Reconcile Planck-scale discreteness and the Lorentz-Fitzgerald contraction. *Physical Review D*, 67:064019, 2003.
- [Smolin, 2000] L. Smolin. *Three Roads to Quantum Gravity*. Oxford University Press, Oxford, 2000.
- [Smolin, 2002] L. Smolin. Matrix models as non-local hidden variables theories, 2002. hep-th/0201031.
- [Smolin, 2003] L. Smolin. How far are we from the quantum theory of gravity?, 2003 hep-th/0303185.
- [Sorkin, 1983] R. Sorkin. Posets as lattice topologies. In B. Bertotti, F. de Felice, A. Pascolini (eds.), *General Relativity and Gravitation: Proceedings of the GR10 Conference*, volume I, page 635. Consiglio Nazionale Delle Ricerche, Rome, 1983.
- [Sorkin, 2002] R. Sorkin. An example relevant to the Kretschmann-Einstein debate. *Modern Physics Letters A*, 17:695–700, 2002.
- [Stachel, 1999a] J. Stachel. Early history of quantum gravity. In B. Iyer and B. Bhawal (eds.), *Black Holes, Gravitational radiation and the Universe*. Kluwer Academic Publisher, Netherlands, 1999a.
- [Stachel, 1999b] J. Stachel. Early history of quantum gravity (1916-1940), 1999b. Presented at the HGR5, Notre Dame, July 1999.
- [Stelle, 1977] K. Stelle. Renormalization of higher derivatives quantum gravity. *Physical Review*, D16:953–969, 1977.
- [Strominger and Vafa, 1996] A. Strominger and G. Vafa. Microscopic origin of the Bekenstein-Hawking entropy. *Physics Letter*, B379:99–104, 1996.
- [’t Hooft, 1973] G. ’t Hooft. An algorithm for the poles at dimension 4 in the dimensional regularization procedure. *Nuclear Physics B*, 62:444–460, 1973.
- [’t Hooft, 1996] G. ’t Hooft. The scattering matrix approach for the quantum black hole, an overview. *International Journal of Modern Physics A*, 11:4623–4688, 1996. gr-qc/9607022.
- [’t Hooft and Veltman, 1974] G. ’t Hooft and M. Veltman. One loop divergencies in the theory of gravitation. *Annales de l’Institut Henri Poincare*, 20:69–74, 1974.
- [Unruh, 1976] W. Unruh. Notes on black hole evaporation. *Physical Review*, D14:870–892, 1976.
- [van Nieuwenhuizen, 1981] P. van Nieuwenhuizen. Supergravity. *Physics Reports*, 68:189–398, 1981.
- [Wald, 1994] R. Wald. *Quantum Field Theory on Curved Spacetime and Black Hole Thermodynamics*. University of Chicago Press, Chicago, 1994.
- [Weinberg, 1979] S. Weinberg. Ultraviolet divergences in quantum theories of gravitation. In S. Hawking and W. Israel (eds.), *General Relativity: An Einstein Centenary Survey*. Cambridge University Press, Cambridge, 1979.
- [Wheeler, 1968] J. Wheeler. Superspace and the nature of quantum geometrodynamics. In C. DeWitt and J. Wheeler (eds.), *Batelle Rencontres, 1967*, volume 242 of *Lectures in Mathematics and Physics*. Benjamin, New York, 1968.
- [Williams, 1997] R. Williams. Recent progress in Regge calculus. *Nuclear Physics Proceedings Suppl.*, 57:73–81, 1997. gr-qc/9702006.
- [Williams and Tuckey, 1992] R. Williams and P. Tuckey. Regge calculus: A bibliography and brief review. *Classical and Quantum Gravity*, 9:1409–1422, 1992.

This page intentionally left blank

SYMMETRIES AND INVARIANCES IN CLASSICAL PHYSICS

Katherine Brading and Elena Castellani

1 INTRODUCTION

The term ‘symmetry’ comes with a variety of ancient connotations, including beauty, harmony, correspondence between parts, balance, equality, proportion, and regularity. These senses of the term are clearly related to one another; the concept of symmetry used in modern physics arose out of this family of ideas. We are familiar with the approximate symmetries of physical objects that we find around us — the bilateral symmetry of the human face, the rotational symmetry of a snowflake turned through 60° , and so forth. We may define a symmetry of a given geometric figure as the invariance of that figure when equal component parts are exchanged under a specified operation (such as rotation). The development of the algebraic concept of a group, in the nineteenth century, allowed a generalization and refinement of this idea; a precise mathematical notion of symmetry emerged which was applicable not just to physical objects and geometrical figures, but also to mathematical equations — and thus, to what is of particular interest to us, the laws of physics expressed as mathematical equations. The group theoretical notion of symmetry is the notion of invariance under a specified group of transformations. ‘Invariance’ is a mathematical term: something is invariant when it is left unaltered by a given transformation. This mathematical notion is used to express the notion of physical symmetry that we are interested in, i.e. invariance under a group of transformations. This is the concept of symmetry that has proved so successful in modern science, and the one that will concern us in what follows.

We begin in Section 2 with the distinction between symmetries of objects and of laws, and that between symmetry principles and symmetry arguments. This section includes a discussion of Curie’s principle. Section 3 discusses the important connection between symmetries, as studied in physics, and the mathematical techniques of group theory. We offer a brief history of how group theory was applied first to geometry and then to physics in the course of the nineteenth century, precluding to the central importance acquired by group theoretical techniques in contemporary physics. With these considerations in mind, Section 4 offers an account of what is meant by symmetry in physics, and a taxonomy of the different

types of symmetry that are found within physics. In Section 5 we discuss some applications of symmetries in classical physics, beginning with transformation theory in classical mechanics, and then turning to Einstein's Special and General Theories of Relativity (see Section 6). We focus on the roles and meaning of symmetries in these theories, and this leads into the discussion of Noether's theorems in Section 7. Finally, in Section 8, we offer some concluding remarks concerning the place, role and interpretation of symmetries in classical physics. Note that our emphasis is resolutely on the classical. For the power and significance of symmetry in quantum physics, we refer the reader to other chapters of this volume, such as Dickson (ch. 4, Section 3.3), Landsman (ch. 5, Section 4.1), t'Hooft (ch. 7) and Halvorson (ch. 8, Section 5.2).¹

2 SYMMETRIES OF OBJECTS AND OF LAWS

That we must distinguish between symmetries of objects versus symmetries of laws can be seen as follows. It is one thing to ask about the geometric symmetries of certain objects — such as the 60° rotational symmetry of a snowflake and the approximate bilateral symmetry of the human face mentioned above — and the *asymmetries* of objects — such as the failure of a chair to be rotationally symmetric. It is another thing to ask about the symmetries of the laws governing the time-evolution of those objects: we can apply the laws of mechanics to the evolution of our chair, considered as an isolated system, and these laws are rotationally invariant (they do not pick out a preferred orientation in space) even though the chair itself is not. Re-phrasing the same point, we should distinguish between symmetries of *states* or *solutions*, versus symmetries of laws. Having distinguished these two types of symmetry we can, of course, go on to ask about the relationship between them: see, for example, current discussions of Curie's principle, referred to in Section 2.2, below.

2.1 *Symmetry principles and symmetry arguments*

It is also important to distinguish between *symmetry principles* and *symmetry arguments*. The application of symmetry principles to laws was of central importance to physics in the twentieth century, as we shall see below in the context of Einstein's Special and General Theories of Relativity. Requiring that the laws — whatever their precise form might be — satisfy certain symmetry properties, became a central methodological tool of theoretical physicists in the process of arriving at the detailed form of various laws.

Symmetry arguments, on the other hand, involve drawing specific consequences with regard to particular phenomena on the basis of their symmetry properties. This type of use of symmetry has a long history; examples include Anaximander's argument for the immobility of the Earth, Archimedes's equilibrium law for the

¹Further discussion can be found in Brading and Castellani [2003].

balance, and the case of *Buridan's ass*.² In each case the associated argument can be understood as an example of the application of the Leibnizean Principle of Sufficient Reason (PSR): if there is no sufficient reason for one thing to happen instead of another, then nothing happens (i.e. the initial situation does not change). There is something more that the above cases have in common: in each of them PSR is applied on the grounds that the initial situation has a certain *symmetry*.³ The symmetry of the initial situation implies the complete equivalence between the offered alternatives. If the alternatives are completely equivalent, then there is no sufficient reason for choosing between them and the initial situation remains unchanged. Arguments of this kind most frequently take the following form: a situation with a certain symmetry evolves in such a way that, in the absence of an asymmetric cause, the initial symmetry is preserved. In other words, a breaking of the initial symmetry cannot happen without a reason: an asymmetry cannot originate spontaneously. This style of argumentation is also to be found in recent discussions of 'Curie's principle', the principle to which we now turn.

2.2 Curie's principle

Pierre Curie (1859-1906) was led to reflect on the question of the relationship between physical properties and symmetry properties of a physical system by his studies on the thermal, electric and magnetic properties of crystals, since these properties were directly related to the structure, and hence the symmetry, of the crystals studied. More precisely, the question he addressed was the following: in a given physical medium (for example, a crystalline medium) having specified symmetry properties, which physical phenomena (for example, which electric and magnetic phenomena) are allowed to happen? His conclusions, systematically presented in his 1894 work 'Sur la symétrie dans les phénomènes physiques', can be summarized as follows:⁴

- (a₁) When certain causes produce certain effects, the symmetry elements of the causes must be found in their effects.
- (a₂) When certain effects show a certain dissymmetry, this dissymmetry must be found in the causes which gave rise to them.⁵
- (a₃) In practice, the converses of these two propositions are not true, i.e., the effects can be more symmetric than their causes.

²For a discussion of these examples, see [Brading and Castellani, 2003, ch. 1, Section 2.2]).

³In the first case rotational symmetry, in the second and third bilateral symmetry.

⁴For an English translation of Curie's paper, see [Curie, 1981]; some aspects of the translation are misleading.

⁵Curie uses the term *dissymmetry* in his paper, as was current at his time. The sense is the same of that of *symmetry breaking* in modern terminology, which is today often identified with the sense of *asymmetry*. To be more precise one should distinguish between the result of a symmetry-breaking process (*broken symmetry*), the absence of *one* of the possible symmetries compatible with the situation considered (*dissymmetry*, as it was called in the nineteenth century literature, notably by Louis Pasteur in his works on molecular dissymmetry), and the absence of *all* the possible symmetries compatible with the situation considered (*asymmetry*).

- (b) A phenomenon may exist in a medium having the same characteristic symmetry or the symmetry of a subgroup of its characteristic symmetry. In other words, certain elements of symmetry can co-exist with certain phenomena, but they are not necessary. What is necessary, is that certain elements of symmetry do not exist. Dissymmetry is what creates the phenomenon.

Conclusion (a₁) is what is usually called *Curie's principle* in the literature. Conclusion (a₂) is logically equivalent to (a₁); the claim is that symmetries are necessarily transferred from cause to effect, while dissymmetries are not. Conclusion (a₃) clarifies this claim, emphasizing that since dissymmetries need not be transferred from cause to effect, the effect may be more symmetric than the cause.⁶ Conclusion (b) invokes a distinction found in all of Curie's examples, between the 'medium' and the 'phenomena'. We have a medium with known symmetry properties, and Curie's principle concerns the relationship between the phenomena that can occur in the medium and the symmetry properties — or rather, '*dissymmetry*' properties — of the medium. Conclusion (b) shows that Curie recognized the important function played by the concept of dissymmetry — of *broken symmetries* in current terminology — in physics.

In order for Curie's principle to be applicable, various conditions need to be satisfied: the cause and effect must be well-defined, the causal connection between them must hold good, and the symmetries of both the cause and the effect must also be well-defined (this involves both the physical and the geometrical properties of the physical systems considered). Curie's principle then furnishes a necessary condition for given phenomena to happen: only those phenomena can happen that are compatible with the symmetry conditions stated by the principle. Curie's principle has thus an important methodological function: on the one hand, it furnishes a kind of selection rule (given an initial situation with a specified symmetry, only certain phenomena are allowed to happen); on the other hand, it offers a falsification criterion for physical theories (a violation of Curie's principle may indicate that something is wrong in the physical description).

Such applications of Curie's principle depend, of course, on our accepting its truth, and this is something that has been questioned in the literature, especially in relation to spontaneous symmetry breaking. Different proposals have been offered for justifying the principle. Curie himself seems to have regarded it as a form of causality principle, and the question in the recent literature has been whether the principle can be demonstrated from premises that include a definition of "cause" and "effect". In this direction it has become current of late to understand the principle as following from the invariance properties of *deterministic* physical laws. The seminal paper for this approach is [Chalmers, 1970], which introduces the formulation of Curie's principle in terms of the relationship between

⁶Note that for some authors conclusion (b) is a principle on its own. Radicati (1987) goes further, describing conclusions (a₁), (a₂) and (b) as three different principles: Curie's first, second and third principle, respectively.

the symmetries of earlier and later states of a system, and the laws connecting these states. This “received view” can be criticized for offering a reformulation that is significantly different from Curie’s intentions (so that the label ‘Curie’s principle’ is a misnomer), and for resting on an assumption that may undermine the interest and importance of the view, as we discuss in the following brief remarks.⁷

The received view, by concerning itself with temporally ordered cause and effect pairs (or states of systems), offers a diachronic or dynamic analysis. In fact, Curie himself focusses on synchronic or static situations, concerning the compatibility of different phenomena occurring at the same time, rather than the evolution of one state of a system into another state. In other words, the ‘cause—effect’ terminology used by Curie is *not* intended to indicate a temporal ordering of phenomena being considered. This is clear from his examples, and also from the fact that discussion of the laws — so central to the diachronic version — is absent from Curie’s own analysis. That the diachronic version has come to have the label ‘Curie’s principle’ therefore misrepresents Curie’s original principle and his discussion of that principle.

Is the diachronic version interesting and important, nevertheless? The account can be understood as an application of PSR in which we pay careful attention to whether the laws provide a “sufficient reason” for a symmetry to be broken as a system evolves from its initial to final state by means of those laws. The reformulation of the diachronic version by Earman [2004] has the strong merit of being precise, and thereby enabling a proof that if the initial state possesses a given symmetry, and the laws deterministically preserve that symmetry, then the final state will also possess that symmetry. However, things are not so simple as they might seem because the proof takes a state with a *given* symmetry. Specifying the symmetries of a state requires, in general, recourse to a background structure — such as space or spacetime, or the space of solutions. In some cases, the required structure may seem trivial or minimal, but nevertheless the dynamics of the system will not be independent of this structure (consider the examples of the spatial or spatiotemporal structure or, more strongly still, the space of solutions). This has the consequence that, in general, the structures on which the symmetries of a state and the symmetries of the dynamics depend are not independent of one another, and any appearance to the contrary in the “proof” needs to be handled with caution. Indeed, we think that answering the question of whether the diachronic version is interesting and important depends in part upon investigating this lack of independence and the role it plays in the proof, something which has yet to be provided in the literature on the diachronic version of ‘Curie’s principle’.

Both Curie’s original version of his principle, and the diachronic version, begin with the symmetries of states of physical systems. In contemporary physics, focus has shifted to symmetries of laws, and the significant connection between symmetries of physical systems and symmetries of laws has to do not with symmetries

⁷For detailed discussion see [Brading and Castellani, 2006]. The “received view” that we attribute first to Chalmers is developed in [Ismael, 1997] and [Earman, 2004]. See also [Earman, this vol., ch. 14, Section 2.3].

of *states* of those systems, but with symmetries of *ensembles of solutions*.⁸ The symmetries of a dynamical equation are not, in general, the symmetries of the individual solutions (let alone states), but rather the symmetries of the whole set of solutions, in the sense that a symmetry of a dynamical equation transforms a given solution into another solution. Considering this relationship between laws and solutions leads to an alternative version of Curie's principle, which we propose here.⁹ As with the diachronic version of Curie's principle, our proposal departs from Curie's original proposal, but our contention is that it remains true to the main motivation behind Curie's original investigation. In this version we seek to unite two things:

1. We understand Curie's motivating question to be 'which phenomena are physically possible?', and his suggestion to be that we can use symmetries as a guide towards answering this question; and
2. We go beyond Curie in making use of symmetries of laws, something about which he said nothing, but which has become a central concern in contemporary physics.

Combining these two ingredients, a "modern" version of Curie's principle would then simply state that *the symmetries of the law (equation) are to be found in the ensemble of its solutions*. This version expresses Curie's basic idea — that "symmetry does not get lost (without a reason)" — in virtue of the fact that the symmetry of the law is to be found in the ensemble of solutions. The fact that this is how we define the relationship between symmetries and laws does not render it empty of significance with respect to Curie's motivating question. On the contrary, the point is that we can *use* the symmetries of the law as a guide to finding solutions, i.e. to determining which phenomena are physically possible, when not all the solutions are known. We can ask, following Curie, 'What phenomena are possible?', and we can use the connection between the symmetries of the law and the symmetries of the ensemble of solutions as a guide to finding the physically possible phenomena. Thus, what is on the one hand a definitional statement (that the symmetries of the law (equation) are to be found in the ensemble of the solutions) comes on the other hand to have epistemic bite when we don't know all the solutions. This, we believe, is true to Curie's motivating question, as expressed in item (1), above.

⁸By "solution" here we mean a temporally extended history of a system, the "state" of a system being a "solution at an instant".

⁹Notice that this version does not involve the temporal evolution from cause to effect (as in the diachronic version), nor is it restricted to a state of a system at a given instant or during a certain temporal period (as in the synchronic version); rather, it concerns the structure of an ensemble of solutions, considered as a whole.

3 SYMMETRY AND GROUP THEORY: EARLY HISTORY

Group theory is the powerful mathematical tool by means of which the symmetry properties of theories are studied. In this section, we begin with the definition of a group, and outline the origins of this notion in the mathematics of algebraic equations. We then turn our attention to the manner in which group theory was applied first to geometry and then to physics in the course of the nineteenth century.

3.1 *The introduction of the group concept and the first developments of group theory*

A *group* is a family \mathcal{G} of elements g_1, g_2, g_3, \dots for which there is defined a multiplication that assigns to every two elements g_i and g_j of the group a third element (their product) $g_k = g_i g_j \in \mathcal{G}$, in such a way that the following requirements hold:¹⁰

- $(g_i g_j) g_k = g_i (g_j g_k)$ for all $g_i, g_j, g_k \in \mathcal{G}$ (associativity of the product);
- there exists an *identity element* $e \in \mathcal{G}$ such that $g_i e = e g_i$ for all $g_i \in \mathcal{G}$;
- for all $g_i \in \mathcal{G}$ there exists an *inverse element* $g_i^{-1} \in \mathcal{G}$ such that $g_i g_i^{-1} = g_i^{-1} g_i = e$.

The concept of a group was introduced by Évariste Galois in the short time he was able to contribute to mathematics (born in 1811, he died as a result of a duel in 1832) in connection with the question of the resolution of equations by radicals.¹¹ The resolvent formulas for cubic and quartic equations were found by the mathematicians of the Renaissance,¹² while the existence of a formula for solving the general equations of the fifth and higher degrees by radicals remained an open question for a long time, stimulating developments in algebra. In particular, the studies of the second half of the eighteenth century focussed on the role played in the solution of equations by functions invariant under permutation of the roots, so giving rise to the theory of permutations. In J. L. Lagrange's *Réflexions sur la résolution algébrique des équations*, the most influential text on the subject, some fundamental results of permutation theory were obtained.¹³ Lagrange's text

¹⁰The concept of a group can be weakened by relaxing these conditions (for example, dropping the inverse requirement leads to the concept of a monoid, and retaining only associativity leads to the concept of a semigroup). The question that then arises is whether the full group structure, or some weaker structure, is related to the symmetry properties of a given theory.

¹¹That is, in terms of a finite number of algebraic operations — addition, subtraction, multiplication, division, raising to a power and extracting roots — on the coefficients of the equations

¹²The resolvent formula for a quadratic equation was known since Babylonian times. A historical survey on the question of the existence of resolvent formulas for algebraic equations is in [Yaglom, 1988, 3 f.].

¹³Among other results, the so-called Lagrange's theorem which states — in modern terminology — that the order of a subgroup of a finite group is a divisor of the order of the group.

served as a basis for successive algebraic developments, from P. Ruffini's first proof in 1799 of the impossibility of solving n th degree equations in radicals for $n \geq 5$, to the seminal works by A. L. Cauchy, N. H. Abel, and finally Galois.¹⁴

Galois's works¹⁵ marked a turning point, providing answers to the open questions in the solutions of equations by using new methods and algebraic notions, first of all the notion of a *group*. This notion was introduced by Galois in relation to the properties of the set of permutations of the roots of equations (the permutations constituting what he named a 'group'), together with other basic notions of group theory such as "subgroup", "normal subgroup", and "simple group".¹⁶ By characterizing an equation in terms of its "degree of symmetry", determined by the permutation group of the roots preserving their algebraic relations (later known as the *Galois group* of the equation), Galois could transform the problem of the resolution of equations into that of studying the properties of the permutation groups involved. In this way he obtained, among other things, the necessary and sufficient conditions for solving equations by radicals.

Galois's achievements in group theory, first brought to publication by Joseph Liouville in 1846, were collected and expanded in Camille Jordan's 1870 *Traité des substitutions et des équations algébriques*. Jordan's *Treatise*, the first systematic textbook on group theory, had a decisive influence on the application of this new theory, including its application to other domains of mathematical science, such as geometry and mathematical physics.

3.2 Applications of group theory: the contributions of Klein and Lie

Projective geometry, the theory of invariants and group theory: Klein and Lie's starting point

In the same year as the appearance of Jordan's *Treatise*, Sophus Lie and Felix Klein, two young mathematicians who were to become the key figures in extending the domain of application of group theory, moved for a period from Berlin to Paris to enter into contact with the French school of mathematics. Lie and Klein had just written a joint paper investigating the properties of some curves in terms of the groups of projective transformations leaving them invariant. In fact they were drawn to Paris mostly by their interest in *projective geometry*, the science founded by J. V. Poncelet to study the properties of figures preserved under central projections. Projective geometry had become, at the time, a particularly fruitful research field for the combination of algebraic and geometrical methods based on the notion of invariance. The *theory of invariants* was itself a flourishing branch

¹⁴Cauchy (1789—1857) generalized Ruffini's results in 1815; Abel (1802—1829) published in 1824 a proof of the impossibility of solving the quintic equation by radicals and in 1826 the paper *Démonstrations de l'impossibilité de la résolution algébrique des équations générales qui passent le quatrième degré*.

¹⁵A few "m'emoires" submitted to the Académie des Sciences, three brief papers published in 1830 in the Académie's 'Bulletin', and some letters, among which is the last one written to his friend Auguste Chevalier in the night before the fatal duel.

¹⁶See [Yaglom, 1988, 9 f.], for details.

of mathematics, centered on the systematic study of the invariants of “algebraic forms”. Using the theory of invariants, the English algebraist Arthur Cayley¹⁷ had recently clarified the relationship between Euclidean and projective geometry, showing the former to be a special case of the latter. Before leaving for France, Klein had tried to extend Cayley’s results, based on the possibility of defining a distance (a “metric”) in terms of a quadratic form defined on the projective space, to the case of non-Euclidean geometries.

While in Paris Lie and Klein became acquainted not only with Jordan, but also with the expert in differential geometry Gaston Darboux, who stimulated their interest in the relations between differential geometry and projective geometry.

Klein’s Erlangen Program

The question of the relations between the different contemporary geometrical systems particularly interested Klein. He aimed at obtaining a unifying foundational principle for the various branches into which geometry had apparently recently separated. In this respect, he fruitfully combined (a) the application of the theory of invariants to the study of geometrical properties, with (b) his and Lie’s idea of applying algebraic group theory to treating also geometrical transformations. The new group theoretical conception of a geometrical theory which resulted was announced in his famous *Erlangen Program*, as it became known following the inaugural lecture entitled ‘Comparative Considerations on Recent Geometrical Research’¹⁸ that the 23-year-old Klein delivered when entering, in 1872, as a professor on the staff of the University of Erlangen. Guided by the idea that geometry is in the end a unity, Klein’s solution to the problem posed by the existence of different geometries was to propose a general characterization of a geometrical theory by using the notion of invariance under a transformation group (i.e., the notion of symmetry). According to his characterization a geometry is defined, with respect to a given domain (the plane, the space, or a given “manifold”) and a group of transformations acting on it, as the science studying the invariants under the transformations of the group. Each specific geometry is thus determined by the characterizing symmetry group (for example, planar Euclidean geometry is determined by the group of affine transformations acting on the plane), and the interrelations between geometries can be described by the relations between the corresponding groups (for example, the equivalence of two geometries amounts to the isomorphism between the corresponding groups).

With Klein’s definition of a geometry, geometrical and symmetrical properties become very close: the symmetry of a figure, which is defined in a given “space”¹⁹ the “geometrical” properties of which are preserved by the transformations of a group \mathcal{G} , is determined by the subgroup of \mathcal{G} leaving the figure invariant. The new

¹⁷Cayley was one of the three members of the ‘invariant trio’, as the French mathematician Hermite dubbed them, the other two being James Sylvester, inventor of most of the terminology of the theory including the word ‘invariant’, and George Salmon.

¹⁸‘Vergleichende Betrachtungen über neuere geometrische Vorschungen’.

¹⁹A set of points endowed with a structure.

group theoretical techniques prompted a transition from an inductive approach (familiar from the nineteenth century classifications of crystalline forms in terms of their visible — and striking — symmetry properties)²⁰ to a more abstract and deductive approach. This is the procedure formulated in Weyl’s classic book on symmetry (*Symmetry*, [1952]) as follows:

Whenever you have to do with a structure-endowed entity Σ try to determine its group of automorphisms, the group of those element-wise transformations which leave all structural relations undisturbed. [...] After that you may start to investigate symmetric configurations of elements, i.e. configurations which are invariant under a certain subgroup of the group of all automorphisms [Weyl, 1952, 144, emphasis in original].

In this way, the symmetry classifications could be extended to figures in “spaces” different from the plane and space of common experience.

Klein himself contributed to the classification of symmetry groups of figures with his works on discrete groups; in particular, he studied the transformation groups related to the symmetry properties of regular polyhedra, which proved to be useful in the solution of algebraic equations by radicals.

Lie’s theory of continuous groups

After 1872, while Klein was concerned especially with *discrete* transformation groups, Lie devoted all his research work to building the *theory of continuous transformation groups*, the results of which were systematically collected in his three-volume *Theorie der Transformationsgruppen* [I: 1888, II: 1890, and III: 1893], written with the collaboration of F. Engel. Lie’s interest in continuous groups arose in relation to the theory of differential equations, which he took to be ‘the most important discipline in modern mathematics’. By the time he was in Paris, Lie had begun to study the theory of first-order partial differential equations, a theory of particular interest because of the central role it played in the formulation given by W. R. Hamilton and C. G. Jacobi to mechanics.²¹ His project was to extend to the case of differential equations Galois’s method for solving algebraic equations: that is, using the knowledge of the ‘Galois group’ of an equation (the symmetry group formed by the transformations taking solutions into solutions) so as to solve it or reduce it to a simpler equation. Thus Lie’s guiding idea was that continuous transformation groups could, in the solution of differential equations, play a role analogous to that of the permutation groups used by Galois in the case of algebraic equations.

²⁰A classic textbook in this respect is [Shubnikov and Koptsik, 1974]. See also Section 8, below.

²¹For details on classical mechanics we refer the reader to [Butterfield, this vol., ch. 1] and the references therein.

Lie had already considered continuous groups of transformations in some earlier geometrical works. In his studies with Klein on special kinds of curves (called by them ‘W-curves’), he had examined transformations that were continuously related in the sense that they were all generated by repeating an *infinitesimal transformation*.²² The relevance of infinitesimal transformations to continuous groups of transformations was to become a central point in his studies of *contact transformations*, so called because they preserved the contact or tangency of surfaces. Lie had started to investigate contact transformations in association with geometrical reciprocities implied in his “line-to-sphere mapping”, a mapping between a line geometry and a sphere geometry that he had discovered while in Paris.²³ When he turned to considering first-order partial differential equations, he soon realized that they admitted contact transformations as symmetry transformations (i.e. transformations taking solutions into solutions). Thus contact transformations could form the “Galois group” of first-order partial differential equations. This motivated him to develop the invariant theory of contact transformations, which represented the first step of his general theory of continuous groups.

Lie’s crucial result, allowing him to pursue his program, was the discovery that to each continuous transformation group could be assigned what is today called its *Lie algebra*. Lie showed that the infinitesimal generators of a continuous transformation group obey a linearized version of the group law, involving the commutator bracket (or Lie bracket); this linearized law then represents the structure of the algebra. In short (and in modern terminology): we describe the elements (transformations) of a continuous group (now called a *Lie group*)²⁴ as functions of a certain number r of continuous parameters a_l ($l = 1, 2, \dots, r$). And these group elements can be written in terms of a corresponding number r of infinitesimal operators X_l , the *generators* of the group, which satisfy the “multiplication law” represented by the Lie brackets

$$[X_s, X_t] = c_{st}^q X_q,$$

so forming what is called the *Lie algebra* of the group. The coefficients c_{st}^q are constants characterizing the structure of the group and are called the *structure constants* of the group.²⁵

Thanks to this sort of result, the study and classification of continuous groups could be conducted in terms of the corresponding Lie algebras. This proved to be extremely fruitful in the successive developments, not only algebraic and geometrical, but also physical. With regard first of all to the physics of Lie’s time, Lie had arrived at the correspondence between continuous groups and Lie algebras by reinterpreting, in the light of his program for solving differential equations, the

²²See on this part [Hawkins, 2000, Section 1.2]. According to Hawkins (p. 15), with the works of Lie and Klein on W-curves ‘for the first time not only is a continuous group the starting point for an investigation, but also for the first time in print we have the idea that infinitesimal transformations are a characteristic and useful feature of continuous systems of transformations’.

²³See [Hawkins, 2000, Section 1.4].

²⁴For a precise definition of this and other terms in this paragraph, see [Butterfield, this vol., ch. 1, Section 3].

²⁵For more details, see [Butterfield, this vol., ch. 1, Sections 3.2 and 3.4].

results obtained by Poisson and especially Jacobi about the integration of first-order partial differential equations arising in mechanics.²⁶ His achievements were thus of great relevance to the solution of the dynamical problems discussed by his contemporaries.

But it is in twentieth century physics, with the works of such figures as Hermann Weyl, Emmy Noether and Eugene Wigner (just to recall the central figures who first contributed to the applications of Lie's theory to modern physics), that the theory of Lie groups and Lie algebras acquired a fundamental role in the description of physical phenomena. Today, the applications of the theory that originated from Lie's works include the whole of theoretical physics, of both the large and the small: classical and quantum mechanics, relativity theories, quantum field theory, and string theory.²⁷

4 WHAT ARE SYMMETRIES IN PHYSICS? DEFINITIONS AND VARIETIES

4.1 What is meant by 'symmetry' in physics

We can understand intuitively the generalization of the scientific notion of symmetry from physical or geometric objects to laws, as follows. We write down our law as a mathematical equation, and appearing in this equation will be various mathematical objects and operators. For a particular group of transformations, these objects and operators transform according to rules that may be fixed either by the mathematical nature of the object or operator concerned, or (where the mathematics does not fix the transformation rules) by our specification. If the "form" of the equation is preserved when we transform each of the objects and operators appearing in our equation by any element of the group, then we say that the group is a *symmetry group* of the equation.

More precisely, what we mean by the symmetry transformations of the laws in physics can be formulated in either of the following ways, which are equivalent in the sense that they pick out the same set of transformations:

- (1) Transformations, applied to the independent and dependent variables of the theory in question, that leave the form of the laws unchanged.
- (2) Transformations that map solutions into solutions.

Symmetry transformations may be viewed either *actively* or *passively*. From the passive point of view we re-describe *the same physical evolution* in two different coordinate systems.²⁸ That is, we transform the independent and dependent variables, as in (1). If the description in the original set of coordinates is a solution

²⁶For details see [Hawkins, 2000, Section 2.5].

²⁷In this volume, see especially t'Hooft (ch. 7), Dickson (ch. 4), and Belot (ch. 2).

²⁸By 'coordinates' here we are referring to generalized coordinates; in general, one coordinate for each degree of freedom of the system.

of given equations, then the new description in the new set of coordinates is a solution of *the same equations*. (If the transformation is *not* a symmetry transformation, then the new description in the new coordinates will not, in general, be a solution of the same equations, but rather of *different* equations.) The mapping of one solution into another solution of the same equations, by means of a symmetry transformation, leads to the *active* interpretation of such transformations. On this interpretation, the two solutions are viewed as *different physical evolutions* described in *the same coordinate system*. Thus, formulation (2) lends itself naturally to an active interpretation.

The ‘form of the law’ in (1) means the functional form of the law, expressed in terms of the independent and dependent variables. A transformation of those variables will, in general, lead to an expression whose functional form differs from that of the original expression (x goes to x^2 , for example). At this point it will be helpful to say a few words about “invariance” and “covariance”. Let the reader beware that there is no unanimity over how these terms are used in discussing the laws of physics, especially in the philosophy of physics literature. Often, the term ‘invariant’ is reserved for objects, and ‘covariant’ is used for equations or laws. However, this is a product of a more fundamental distinction, which when understood correctly allows for the application of the notion of invariance to laws as well.

We think that the discussion of Ohanian and Ruffini [1994, Section 7.1] is very useful, and that it nicely distills much of the best of what can be found in the literature, both in physics and in philosophy of physics. The upshot is as follows. We may say that an equation is *covariant* under a given transformation when its form is left unchanged by that transformation. This is the notion at work in Definition 1. In a way, it is rather weak: given an equation that is not covariant under a given transformation, we can always re-write it so that it becomes covariant. On the other hand, this re-writing may involve the introduction of new functions of the variables, and it is the physical interpretation of these new quantities that allows covariance to gain physical significance. We will have more to say about this for the specific case of *general covariance* and Einstein’s General Theory of Relativity in Section 6.3 below.

Invariance of an equation, as characterized by Ohanian and Ruffini, is a stronger requirement than covariance. Not only should the form of the equation remain the same, but so too should the values of any non-dynamical quantities, including “constants” such as the speed of light. By “non-dynamical quantities” we mean all those objects which appear in the equations yet which do not themselves satisfy equations of motion. We here enter the muddy waters of how to distinguish between “absolute” and “dynamical” objects, as discussed by Anderson [1967].²⁹

In both cases (covariance and invariance), the associated transformations — when actively construed — take solutions into solutions. When using formulation

²⁹See also Section 6.3, below. One difficulty in tackling the literature on this issue is the variety of uses and meanings attaching to the common terminology of covariance, principle of covariance, invariance, absolute and dynamical objects, and so forth.

(2), it is important to be clear about what is meant by a solution. This does *not* mean a solution-at-an-instant, i.e. an instantaneous state of a system; rather, it means an entire history, i.e. possible time-evolution, of the system in question.³⁰

4.2 Varieties of symmetry

Symmetries in physics come in a number of different varieties, distinguished by such terms as ‘global’ and ‘local’; ‘internal’ and ‘external’; ‘continuous’ and ‘discrete’. In this Section we briefly review this terminology and the associated distinctions.

The most familiar are the global spacetime symmetries, such as the Galilean invariance of Newtonian mechanics, and the Lorentz invariance of the Special Theory of Relativity. Global spacetime symmetries are intended to be valid for all the laws of nature, for all the processes that unfold in the spacetime. Symmetries with this universal character were labelled ‘geometric’ by Wigner (see [1967, especially p. 17]).

This universal character is not shared by some of the symmetries introduced into physics during the twentieth century. Most of these were of an entirely new kind, with no roots in the history of science, and in some cases expressly introduced to describe specific forms of interactions — whence the name ‘dynamical symmetries’ due to Wigner [1967, see especially pp. 15, 17—18, 22—27, 33]).

The various symmetries of modern physics can also be classified according to a second distinction: that between global and local symmetries. The terms ‘global’ and ‘local’ are used in physics, and in philosophy of physics, with a variety of meanings. The distinction intended here is between symmetries that depend on constant parameters (global symmetries) and symmetries that depend on arbitrary smooth functions of space and time (local symmetries). While Lorentz invariance is an example of a global symmetry, the gauge symmetry of classical electromagnetism (an internal symmetry)³¹ and the diffeomorphism invariance in General Relativity (a spacetime symmetry) are examples of local symmetries, since they are parameterized by arbitrary functions of space and time.³² Recalling Wigner’s distinction, Lorentz invariance is a geometric symmetry, applying to all interactions, whereas the gauge symmetry of electromagnetism concerns the electromagnetic interaction specifically and is therefore a dynamical symmetry.

The gauge symmetry of classical electromagnetism is an internal symmetry because the transformations of the vector potential occur in the internal space of the field system, rather than in spacetime. The gauge symmetry of classical electromagnetism can seem to be no more than a mathematical curiosity, specific to this theory; but with the advent of quantum theory the use of internal degrees

³⁰The distinction is important in, for example, our discussion of Curie’s principle, Section 2.2 above.

³¹For more on gauge and internal symmetries, see the following paragraph.

³²We discuss the local symmetry of General Relativity further in Section 6.1 below. See also [Belot, this vol., ch. 2].

of freedom, and the related internal symmetries, became fundamental.³³

The translations, rotations and boosts of the inhomogeneous Lorentz group are all examples of *continuous symmetries*, for which any finite symmetry transformation can be built up of infinitesimal symmetry transformations. In contrast with the continuous symmetries we have the *discrete symmetries* of charge conjugation, parity, and time reversal (CPT), along with permutation invariance. Thus, Newtonian mechanics and classical electrodynamics are invariant under parity (left-right inversion) and under time reversal (roughly: the laws hold for a sequence of states evolved in the backwards time direction just as they hold for the states ordered in the forwards direction). Classical electrodynamics is also invariant under charge conjugation, so long as we correctly implement the associated transformations of the electric and magnetic fields. Finally, there is a sense in which classical statistical mechanics is permutation invariant: the particles postulated are identical to one another, and their permutation takes a solution into a solution. However, the power and significance of the discrete symmetries achieves its full force only in quantum theory.

In Section 8 below, we discuss some of the interpretative issues associated with these different varieties of symmetry in classical physics.

5 SOME APPLICATIONS OF SYMMETRIES IN CLASSICAL PHYSICS

5.1 *Transformation theory in classical mechanics*

As we have seen, Lie's interest in continuous groups arose in relation to his studies of the theory of first-order partial differential equations, which played a central role in the formulation given by Hamilton and Jacobi to mechanics. The *transformation theory of mechanics* based on this formulation is indeed one of the first examples of a systematic exploitation in physics of the invariance properties of dynamical equations. These symmetries are exploited according to the following strategy: the integration of the equations of motion is simplified by transforming — by means of symmetry transformations — the original dynamical system into another system with fewer degrees of freedom.

Historically, the road to the possibility of applying the above 'transformation strategy' to solving dynamical problems was opened by the works of J. L. Lagrange and L. Euler. The Euler-Lagrange analytical formulation of mechanics, grounded in the seminal *Mécanique Analytique* [1788] of Lagrange, expressed the laws of motion in a form which was covariant (cf. Section 4.1) under all coordinate transformations. This meant one could more easily choose coordinates to suit the dynamical problem concerned. In particular, one hoped to find a coor-

³³For interpretative issues associated with gauge symmetry in classical electromagnetism, see Belot [1998]. Gauge symmetries came to prominence with the development of quantum theory. The term 'gauge symmetry' itself stems from Weyl's 1918 theory of gravitation and electromagnetism. For discussions of all these aspects of gauge symmetry, see [Brading and Castellani, 2003].

dinate system containing “cyclic” (a.k.a. “ignorable”) coordinates. The presence of ignorable coordinates amounts to a partial integration of the equations: if all the coordinates are ignorable, the problem is completely and trivially solved. The method was thus to try to find (by applying coordinate transformations leaving the dynamics unchanged) more and more ignorable variables, thus transforming the problem of integrating the equations of motion into a problem of finding suitable coordinate transformations.

The successive developments in the analytical approach to mechanics, from Hamilton’s “canonical equations of motion” to the general transformation theory of these equations (the theory of canonical transformations) obtained by Jacobi, presented many advantages of the “transformation strategy” point of view. For further details we refer the reader to Butterfield’s chapter of this volume, along with classic references such as Lanczos [1949; 1962] and Whittaker [1904; 1989]. Butterfield [this vol., ch. 1], by expounding the theory of symplectic reduction in classical mechanics, thoroughly illustrates the strategy of simplifying a mechanical problem by exploiting a symmetry. This strategy is also the main subject of Butterfield [2006], focussing on how symmetries yield conserved quantities according to Noether’s first theorem (see Section 7, below), and thereby reduce the number of variables that need to be considered in solving a problem.

We end these brief remarks on symmetry and transformation theory in classical mechanics by emphasizing two points.

First, we note that a problem-solving strategy according to which a dynamical problem (equation) is transformed into another equivalent problem (equation) by means of a symmetry might be seen as an example of the application of Curie’s principle in its modern version (see here Section 2.2): by transforming an equation into another equivalent equation using a specific symmetry we may arrive at an equation which we can solve; the solution of the new equation is related to the unknown solution of the old equation by the specific symmetry; that is, we thereby arrive at an equivalent solution.

Second, we emphasize that in all these developments the invariance properties of the dynamical equations, though undoubtedly important, were considered exclusively in an instrumental way. That is, canonical transformations were studied only for the purpose of solving the dynamical problem at hand. The equations were given, and their invariance properties were investigated to help find their solutions. The formulation of Einstein’s Special Theory of Relativity at the beginning of the twentieth century brings an inversion of this way of thinking about the relationship between symmetries and physical laws, as we shall see in the following section.

5.2 Symmetry principles as guides to theory construction

The principle of relativity, as expressed by Einstein in his 1905 paper announcing the Special Theory of Relativity, asserts that

The laws by which the states of physical systems undergo changes are

independent of whether these changes are referred to one or the other of two coordinate systems moving relatively to each other in uniform motion.³⁴

It further turns out that these coordinate systems are to be *inertial* coordinate systems, related to one another by the Lorentz transformations comprising the inhomogeneous Lorentz group.

The principle of relativity thus stated meets the conditions listed above in Section 1 for a *symmetry* principle:

- The Lorentz transformations, applied to the independent and dependent variables of the theory, leave the form of the laws as stated in one inertial system unchanged on transformation to another inertial coordinate system.
- The Lorentz transformations map a solution, given relative to an inertial coordinate system, into another solution.

This principle was explicitly used by Einstein as a guide to theory construction: it is a principle that must be satisfied whatever the final details of the theory.³⁵ Indeed, using just the principle of relativity and the light postulate, Einstein derives various results, including the Lorentz transformations. As noted above, this represents a reversal in the priority that, since the time of Newton, had been given to the relativity principle versus the dynamical laws. Huygens used the relativity principle as a basic postulate from which to derive dynamical results, but in Newton the relativity principle, initially presented in his manuscripts as an independent postulate, is relegated in the *Principia* to a corollary.³⁶ From then until Einstein, the relativity of inertial motion is seen as a consequence of the particular laws under consideration, and something that could turn out to be false once the details of the laws of some particular interaction are known. Similarly for classical physics in general, symmetries — such as spatial translations and rotations — were viewed as properties of the laws that hold as a consequence of those particular laws. With Einstein that changed: symmetries could be postulated prior to details of the laws being known, and used to place restrictions on what laws might be postulated. Thus, symmetries acquired a new status, being postulated independently of the details of the laws, and as a result having strong heuristic power. As Wigner wrote, Einstein's papers on special relativity 'mark the reversal of a trend': after Einstein's works, 'it is now natural to try to derive the laws of nature and to test their validity by means of the laws of invariance, rather than to derive the laws of invariance from what we believe to be the laws of nature' [Wigner, 1967, 5].

³⁴Miller's [1981] translation, p. 395.

³⁵For discussion of the principle/constructive theory distinction in Einstein, see [Brown, 2006, ch. 6] and [Howard, 2007].

³⁶In fact, it does not follow from Newton's three laws of motion — we must further assume the velocity independence of mass and force. See [Barbour, 1989, Section 1.2].

The methodology that had served Einstein well with the Special Theory of Relativity (STR) also had a role in his development of the General Theory (GTR), for which he used various different principles as restrictions on the possible form that the eventual theory might take.³⁷ One of these was, so Einstein maintained, an extension of the principle of relativity found in STR to include coordinate systems that are in accelerated motion relative to one another, implemented by means of the requirement that the equations of his new theory be generally covariant. Einstein was seeking a “Machian” solution to the challenge of Newton’s bucket, which he took to require that there be no preferred reference frames. Thus, in his 1916 review article Einstein wrote that ‘*The laws of physics must be of such a nature that they apply to systems of reference in any kind of motion.* Along this road we arrive at an extension of the postulate of relativity’ (emphasis in original).

The questions of whether or not the principle of general covariance (a) makes any arbitrary smooth coordinate transformation into a symmetry transformation, and (b) is a generalization of the principle of relativity, have been much discussed. The answer to (b) is a definitive ‘no’, but there is less consensus at present about the answer to (a).³⁸ In the following section we take up discussion of (a). Here we close with a few brief remarks concerning (b).

Even if general covariance in GTR is a symmetry principle, it is not an extension of the relativity principle. That is to say, general covariance says nothing about the observational equivalence of distinct reference frames.³⁹ As already noted, the thought that general covariance might provide such a principle was, for Einstein, connected with his attempts to provide a “Machian” resolution to the challenge of Newton’s bucket, and with his principle of equivalence. However, the principle of equivalence *does not* imply the observational equivalence of reference frames in arbitrary states of motion (Einstein never thought that it did), and Einstein eventually realized that GTR does not vindicate a solution to Newton’s bucket that depends only on the relative motion of *matter*.⁴⁰

Whatever the subtleties of whether, and to what extent, general covariance is a symmetry principle, it is clear that it had enormous heuristic power, not just for Einstein in his development of GTR, but also beyond. Think for example Hilbert’s work on the axiomatization of physics (see [Corry, 2004], and references therein), and Weyl’s attempts to construct a unified field theory (see [O’Raifeartaigh, 1997], for an English translation of Weyl’s 1918 paper ‘Gravitation and Electricity’, and see also [Weyl, 1922]). In all these cases, general covariance provided a powerful tool for theory construction. In the following Section we discuss further the

³⁷Primarily the following: the principle of relativity, later (in 1918) distinguished from what Einstein referred to as ‘Mach’s principle’; the principle of equivalence; and the principle of conservation of energy–momentum.

³⁸See for example [Torretti, 1983, 152–4]; Norton [1993], who also discusses the relationship with the principle of equivalence; Anderson [1967].

³⁹For further discussion see, for example, [Norton, 1993] and [Torretti, 1983, Section 5.5].

⁴⁰For a clear and concise discussion, see [Janssen, 2005].

significance and interpretation of general covariance in GTR.⁴¹

6 GENERAL COVARIANCE IN GENERAL RELATIVITY

In the preceding Section we noted the role of the principle of general covariance as a guide to theory construction. In this Section we turn our attention to a number of further issues relating to general covariance in GTR that have received attention in the philosophical literature. We begin with the issue, raised in the preceding section, concerning the status of arbitrary smooth coordinate transformations as symmetry transformations. We then discuss various characteristics associated with general covariance, including those pointed to by Einstein's so-called 'hole argument', before turning to the issue of whether or not general covariance has physical content.⁴² We postpone discussion of Noether's theorems to Section 7, below.

6.1 *General covariance and arbitrary coordinate transformations as symmetry transformations*

Does the principle of general covariance make any arbitrary smooth coordinate transformation into a symmetry transformation? One way to approach this question is to consider active rather than passive transformations (see Section 4.1, above), and to compare the situation in GTR with that in STR.

In STR, a Lorentz transformation — actively construed — picks up the matter fields and redistributes them with respect to the spacetime structure encoded in the metric. The principle of relativity holds for such transformations because the evolution of the matter fields in the two cases (related by the Lorentz transformation) are observationally indistinguishable: no observations, in practice or in principle, could distinguish between the two scenarios. In GTR, active general covariance is implemented by active diffeomorphisms on the spacetime manifold (see [Rovelli, this vol., ch. 12, Section 4.1]). These involve transformations of not just the matter fields, but also the metric field, in which both are redistributed with respect to the spacetime manifold. Once again, the "two cases" are observationally "indistinguishable", but this time the reason generally given is that the "two cases" are in fact just one case.⁴³

Why should we accept that there are two genuinely distinct cases when considering the Lorentz transformations in STR, and only one case for the diffeomorphisms of GTR? One approach would be to claim that a crucial difference between the two is that a Lorentz transformation can be implemented on an effectively isolated *sub-system* of the matter fields, producing an observably distinct scenario in

⁴¹For detailed presentation of the Special and General Theories of Relativity, see [Malament, this vol., ch. 3]. See also [Rovelli, this vol., ch 12, Section 4].

⁴²See also [Belot, this vol., ch. 2].

⁴³See also Section 6.2, penultimate paragraph.

which, nevertheless, the evolution of the sub-system in question is indistinguishable assuming no reference is made to matter fields outside that subsystem. For example, in Galileo's famous ship experiment we consider two observably distinct scenarios — one in which the ship is at rest with respect to the shore, and one in which it moves uniformly with respect to the shore — and we notice that the behaviour of physical systems within the cabin of the ship does not distinguish between the two scenarios.⁴⁴ No analogue of the Galilean ship experiment can be generated for the general covariance of GTR.⁴⁵

The importance of symmetry transformations being implementable to produce observationally distinct scenarios has been emphasized by Kosso [2000]. On this view, the observational significance of symmetry transformations rests on a combination of *two* observations being possible in principle. First, it must be possible to confirm empirically the implementation of the transformation — hence the importance of being able to generate an observationally distinguishable scenario through the transformation of a subsystem. Second, we must be able to observe that the subsequent internal evolution of the subsystem is unaffected. That we cannot meet the first of these requirements for arbitrary smooth coordinate transformations in GTR marks a difference between these and the Lorentz transformations.⁴⁶

On this approach, while the field equations of GTR take the same form for any choice of coordinate system, this is not sufficient for arbitrary coordinate transformations to be symmetries. In addition, the actively construed transformations must have a physical interpretation — we must be transforming one thing with respect to something else. When we perform a diffeomorphism, we get back the same solution, not a new solution, for we are not re-arranging the matter fields with respect to the metric.

We stress that this is only one way to approach the issue of whether general covariance should be understood as a symmetry principle in GTR. A contrasting position may be found in [Anderson, 1967, Section 10-3], who argues that we must understand Einstein as viewing general covariance as a symmetry requirement, and attempts to spell out the conditions under which it can function as such.

⁴⁴This implementation can be only approximate, relying on the degree to which the subsystem in question can be isolated from the “external” matter fields.

⁴⁵One suggestion might be that we perform a transformation T which is the identity outside some region R , and which differs from the identity within that region. This will not achieve the desired result. The two scenarios must have observationally distinct consequences, at least in principle. In the case of Galileo's ship, if we allow the subsystem to interact with other matter once again, we will see that in one case the ship crashes into rocks (for example), while in the other it suffers no such collision. Thus, we have observational distinguishability in principle. The transformation T does not produce a scenario which any future events could enable us to distinguish from the original.

⁴⁶Indeed, this result applies generally to local versus global symmetries. See also [Brading and Brown, 2004].

6.2 Characteristics of generally covariant theories

Any generally covariant theory will possess certain characteristics that are philosophically noteworthy. First, there will be a *prima facie* problem with causality and determinism within the theory, and second, there will be constraints on the specification of the initial data. Einstein recognized aspects of the first characteristic while he was searching for his theory of gravitation, maintaining from 1913 through until the fall of 1915 that his so-called ‘hole argument’ provided grounds for concluding that no generally covariant theory could be physically acceptable.

In the ‘hole argument’, Einstein considers a region of spacetime in which there are no matter fields (the “hole”), and then shows that in a generally covariant theory no amount of data about the values of the matter and gravitational fields outside the hole is sufficient to uniquely determine the values of the gravitational field inside the hole. From this, Einstein concluded that no generally covariant theory could be physically acceptable.⁴⁷

The context to bear in mind here is that Einstein was searching for a theory in which the matter fields plus the field equations would uniquely determine the metric.⁴⁸ In the summer of 1915 Einstein lectured on relativity theory in Göttingen where his audience included David Hilbert. If we assume that Einstein’s presentation included a version of his ‘hole argument’, then we can reasonably infer that Hilbert was quick to reinterpret the issue that the ‘hole argument’ points to, and to present the problem raised for generally covariant theories in terms of whether such theories permit well-posed Cauchy problems.⁴⁹

In the years immediately following the advent of GTR, Hilbert played a central role in spelling out the problems of causality and determinism faced by any generally covariant theory. He pointed out that in any such theory, including GTR, there will be four fewer field equations than there are variables, leading to a mathematical underdetermination in the theory. As Hilbert stressed, the Cauchy problem is not well-posed: given a specification of initial data, the field equations do not determine a unique evolution of the variables.

We can see the connection between the underdetermination problem and general covariance as follows. For the Cauchy problem to be well-posed, we must be able to express the second time derivatives of the metric in terms of the initial data (plus the further spatial derivatives that can be calculated from the initial data). However, if we re-express the 10 (source-free) Einstein field equations $G_{\mu\nu} = 0$ so as to explicitly display all the terms containing the second time derivative of the metric, we see that we have ten equations for six unknowns $g_{ij,00}$, the remaining four second time derivatives $g_{\mu 0,00}$ failing to appear in the equations.⁵⁰ This

⁴⁷For presentation and discussion of the ‘hole argument’, see Norton [1984, 286–291] and [1993, Sections 1-3], Stachel [1993], and Ryckman [2005, Section 2.2.2]. See also [Rovelli, this vol., ch. 12, Section 4.1.1].

⁴⁸For more on Einstein’s (mis)appropriation of Mach’s principle, see [Barbour, 2005].

⁴⁹Brading and Ryckman [2007]; see also [Brown and Brading, 2002, especially Section IV].

⁵⁰See [Adler, Bazin and Schiffer, 1975, ch. 8] for details of the over- and under-determination issues.

is a direct consequence of general covariance: we can always make a coordinate transformation in the neighborhood of the initial data surface such that the metric components and their first derivatives are unchanged, while the second time derivatives $g_{\mu 0,00}$ vanish on that surface. Thus the field equations, which must be valid in all coordinate systems, cannot possibly contain information on the second time derivatives. The initial data do not determine the metric uniquely: there are four arbitrary functions $g_{\mu 0,00}$ that we are free to choose.

Today, it is customary to assert from the outset that solutions of Einstein's field equations differing only in the choice of these four arbitrary functions are physically equivalent.⁵¹ But here we should note that this "gauge freedom" interpretation of general covariance leads to problems of its own.⁵² For example, within this framework the observables of the theory must be "gauge invariant" quantities, but such quantities have (to date) turned out to be far removed from anything "observable" in the operational sense. The gauge freedom interpretation of general covariance is sometimes accompanied by the view that this freedom — and therefore general covariance itself — lacks physical content. We turn to consider this issue in Section 6.3, below.

In our explanation of the underdetermination problem, above, we noted that the Einstein equations provide ten equations for the six unknowns $g_{ij,00}$. The other face of the underdetermination problem is therefore an *overdetermination* problem with respect to the $g_{ij,00}$, and what this means is that there will be constraints on the specification of the data on the initial hypersurface. This is the second characteristic of all generally covariant theories that we mentioned in our opening remarks of the current subsection. Indeed, the presence of constraint equations is a feature shared with other theories with a local symmetry structure, such as electromagnetism. Philosophically, the significance lies in the relationship between the theory and the initial data. In the seventeenth century Descartes wrote a story of a world created in a state of disorder from which, by the ordinary operation of the laws of nature, a world seemingly similar to our own emerged.⁵³ This image of the world emerging from an initial chaos has a long history, of course, but the emergence of order by means of the operation of the laws of nature offered a novel twist to the tale. It involves the separation of initial conditions, which could be anything, from the subsequent law-governed evolution of the cosmos. In modern terms, this is a theory without constraints: the theory determines *which* properties of a system must be specified in order to give adequate initial data, but we are then free to assign whatever *values* we please to these properties; the equations of the theory are used to evolve that data forwards in time. A theory with constraints, by contrast, contains two types of equations: constraint equations that must be satisfied by the initial data, as well as evolution equations.

⁵¹Recall the discussion of Section 6.1, above.

⁵²See [Belot, this vol., ch. 2].

⁵³Written around 1633, *Le Monde* was not published in Descartes's lifetime. For an English translation see Descartes [1998]. The "order out of disorder" story is in the Treatise on Light, chs. 6 and 7. Whether the ordinary operation of the laws of nature was sufficient to bring order out of chaos became a much-disputed issue.

In GTR, four of the ten field equations connect the curvature of the initial data hypersurface with the distribution of mass–energy on that hypersurface, and the remaining six field equations are evolution equations. To sum up, in a theory with constraints, the initial “disorder” cannot be so disordered after all, but must itself satisfy constraints set down by the laws of the theory.

6.3 *Does general covariance have any physical significance?*

As we saw in Section 5, Einstein treated general covariance as a symmetry principle guiding the search that produced his General Theory of Relativity. There is no doubt that general covariance proved a useful heuristic for Einstein, but there remains an ongoing dispute over whether general covariance in fact has any physical significance. The issue was forcefully raised by Kretschmann already in 1917. The thrust of the argument, which continues to reverberate today, is that any theory can be given a generally covariant formulation given sufficient mathematical ingenuity, and therefore the principle of general covariance places no restrictions on the physical content of a theory. Indeed, Norton [2003] begins his discussion of the issue by claiming that this negative view of general covariance has become mainstream, before going on to give an alternative viewpoint (see below).

It seems clear to us that the characteristic features of generally covariant theories discussed above may, in some theories at least (including GTR), be far from trivial, and that the mainstream view — which would indeed render these issues trivial — should be opposed. Those wishing to oppose the mainstream view adopt a two-step general strategy: first, show under what conditions general covariance places a restriction on the physical content of a theory; and second, demonstrate what those implications for physical content consist in. Thus, the general mathematical point that any theory can be put into generally covariant form is conceded, but the implication that general covariance is therefore necessarily physically vacuous is resisted by attention to the manner in which general covariance is implemented in a given theory or class of theories.

For example, Anderson [1967], Ohanian and Ruffini [1994], Norton [2003], and Earman [2006] each attempt to explain under what conditions the purely mathematical feature of general covariance comes to have physical bite.⁵⁴ Anderson distinguishes between the symmetries of a theory (which have physical significance) and the covariance group of the equations (which need not). Anderson is the classic reference for the distinction between “absolute” and “dynamical” objects,⁵⁵

⁵⁴See also [Norton, 1993, especially Section 5], and [Rovelli, this vol., ch. 12, Section 4.1.3].

⁵⁵It has proved difficult to make the distinction between absolute and dynamical objects precise, but the intuitive idea is clear enough. Dynamical objects satisfy field equations and interact with other objects, whereas absolute objects are not affected by the dynamical behaviour of other fields appearing in the theory. For a careful and detailed treatment of Anderson’s approach, and the counter-examples that have been raised, see [Pitts, 2006]. The conclusion of this paper is that Anderson’s intuition can be made sufficiently precise to cope with all counter-examples that have appeared in the literature to date (including one due to Pitts himself), but that there is another example, due to Geroch, that Pitts has been unable to resolve. The debate goes on!

and in this terminology the covariance group of the equations of a theory becomes a symmetry group if and only if the theory contains no absolute objects. Ohanian and Ruffini [1994] appeal to the distinction they make between *invariance* and *covariance* of the equations of a theory.⁵⁶ Covariance, they agree, is a mathematical feature (perhaps simply an artefact of the particular formulation of the theory at hand); but we require not only the covariance of the equations, but also that for any objects (with one or more components) appearing in the theory that are nevertheless independent of the state of matter (such as the speed of light, Planck's constant, etc.), their value should be unchanged by the general coordinate transformations. Norton [2003] emphasizes the role of physical considerations in fixing the content of a theory such that this restricts the formal games that we can play. Earman [2006] begins by taking pains to emphasize the distinction between the 'mere co-ordinate freedom' (associated with arbitrary coordinate transformations, passively construed) and 'the substantive demand that diffeomorphism invariance is a gauge symmetry of the theory at issue'. That is to say, he reminds us that the issue at stake is not our ability to re-write a theory in generally covariant form (it is conceded that this is something we can always do, given sufficient mathematical ingenuity), but the relationship between the physical situations that are related by diffeomorphisms, i.e. by (active) point transformations (see Section 6.1, above). 'Substantive general covariance' holds when diffeomorphically related models of the theory represent different descriptions of the same physical situation. The claim is that GTR satisfies substantive general covariance whereas generally covariant formulations of such theories as STR need not, and the goal is to show that this requirement provides demarcation between theories in which general covariance represents a physically significant property of the theory, and those in which it does not.⁵⁷

Thus, Anderson, Ohanian and Ruffini, Norton, and Earman each seek to add bite to the "merely mathematical" requirement of general covariance by placing conditions on the manner in which it is implemented in the theory. Once these requirements are added, various consequences follow for the content of the theory, such as that the metric be a dynamical object. In each case, the aim is to elevate general covariance as implemented in GTR to a symmetry principle.⁵⁸

Considerations of the significance of general covariance in theories of gravitation led to the formulation of three theorems important for the general interpretation of symmetries in physics. These theorems are due to Emmy Noether and Felix Klein, and will be discussed in the following section.

⁵⁶See Section 4.1.

⁵⁷One important tool for distinguishing genuine 'gauge theories' from those in which the local symmetry in question is merely formal is Noether's second theorem; see Section 7, below.

⁵⁸Brown and Brading [2002] attempt to analyze in more detail, by means of Noether's theorems (see Section 7, below), what additional conditions must be added to general covariance in order to arrive at specific aspects of the content of GTR.

7 NOETHER'S THEOREMS

Any discussion of the significance of symmetries in physics would be incomplete without mention of Noether's theorems. These theorems relate symmetry properties of theories to other important properties, such as conservation laws.

Within physics, the term 'Noether's theorem' is most frequently associated with a connection between global continuous symmetries and conserved quantities. Familiar examples from classical mechanics include the connections between: spatial translations and conservation of linear momentum; spatial rotations and conservation of angular momentum; and time translations and conservation of energy. In fact, this theorem is the first of two theorems presented in her 1918 paper 'Invariante Variationsprobleme'.⁵⁹

Before stating the two theorems, we begin with the following cautionary remark. The connection between *variational symmetries* (connected to the invariance of the action, and in terms of which Noether's theorems are formulated) and *dynamical symmetries* (concerning the dynamical laws, which is the topic of our discussion here) is subtle (see [Olver, 1993, ch. 4]). Noether herself never addressed the connection, and never used the word 'symmetry' in her paper. She discusses integrals mathematically analogous to (but generalizations of) the action integrals of Lagrangian physics, and uses variational techniques and group theory to elicit a pair-wise correspondence between variational symmetries of the integral and a set of identities.

Noether then proves two theorems, the first for the case where the variational symmetry group depends on constant parameters, and the second for the case where the variational symmetry group depends on arbitrary functions of the variables.⁶⁰ In the following statement of her theorems we use the term 'Noether symmetry' to refer to a symmetry of the field equations for which the change in the action arising from the infinitesimal symmetry transformation is at most a surface term. Using the terminology of Section 4.2, the first type of symmetry then corresponds to a *global* dynamical symmetry, and the second to a *local* one. We state the theorems in a form appropriate to Lagrangian field theory; Noether's own statement of the theorems involves no such specialization. For discussion of the first theorem in the context of finite-dimensional classical mechanics see [Butterfield, this vol., ch. 2, Section 2.1.3]. We state the theorems so that we can refer back to them to characterize the conceptual content, but for discussion of the mathematical detail of their derivation and content we refer the reader elsewhere — see especially [Olver, 1993] and [Barbashov and Nesterenko, 1983].

We can state Noether's two theorems, for a Lagrangian density L depending on the fields $\phi_i(x)$ and their first derivatives, as follows.

⁵⁹For an English translation see [Noether, 1971].

⁶⁰See [Brading and Brown, 2007].

Noether’s first theorem

If a continuous group of transformations depending smoothly on ρ constant parameters ω_k ($k = 1, 2, \dots, \rho$) is a Noether symmetry group of the Euler-Lagrange equations associated with a Lagrangian $L(\phi_i, \partial_\mu \phi_i, x^\mu)$, then the following ρ relations are satisfied, one for every parameter on which the symmetry group depends.⁶¹

$$\sum_i E_i^L \xi_i^k = \partial_\mu j_k^\mu. \tag{1}$$

On the left-hand side we have a linear combination of Euler expressions,

$$E_m^L \equiv \frac{\partial L}{\partial \phi_m} - \partial_\mu \left(\frac{\partial L}{\partial \phi_{m,\mu}} \right) \tag{2}$$

where

$$E_m^L = 0 \tag{3}$$

are the Euler-Lagrange equations for the field ϕ_m . (The ξ_i^m depend on the particular symmetry transformations and fields under consideration, and the details are not important for our current purposes.)

On the right-hand side we have the divergence of a current, j_k^μ . When the left-hand side vanishes, the divergence of the current is equal to zero, and this expression can be converted into a conserved quantity subject to certain conditions. Thus, Noether’s first theorem gives us a connection between global symmetries and conserved quantities.⁶²

Noether’s second theorem

If a continuous group of transformations depending smoothly on ρ arbitrary functions of time and space $p_k(x)$ ($k = 1, 2, \dots, \rho$) and their first derivatives is a Noether symmetry group of the Euler-Lagrange equations associated with a Lagrangian $L(\phi_i, \partial_\mu \phi_i, x^\mu)$, then the following ρ relations are satisfied, one for every function on which the symmetry group depends:

$$\sum_i E_i^L a_{ki} = \sum_i \partial_\nu (b_{ki}^\nu E_i^L). \tag{4}$$

The a_{ki} and b_{ki}^ν depend on the particular transformations of the fields in question, and while again the details need not concern us here, we note for use below that while the a_{ki} arise even when the symmetry transformation is a global transformation, the b_{ki}^ν occur only when it is *local*.⁶³ What we have here, essentially,

⁶¹Note that we are using the Einstein summation convention to sum over repeated greek indices.

⁶²This theorem is widely discussed. See especially [Barbashov and Nesterenko, 1983]; [Doughty, 1990]. We refer the reader to [Butterfield, this vol., ch. 2] and [Butterfield, 2006] for further discussion of Noether’s first theorem in the context of finite-dimensional classical mechanics.

⁶³Once again, the reader is referred to [Brading and Brown, 2007] for further details.

is a dependency between the Euler expressions and their first derivatives. This dependency holds as a consequence of the *local* symmetry used in deriving the theorem. In the case when all the fields are dynamical (i.e. satisfy Euler-Lagrange equations) it follows that not all the field equations are independent of one another. This *formal underdetermination* is characteristic of theories with a local symmetry structure.⁶⁴

As Hilbert recognized in the context of generally covariant theories of gravitation, the underdetermination is independent of the specific form of the Lagrangian.⁶⁵ In the case of General Relativity, once we specify the Lagrangian and substitute it into (4), we arrive at the (contracted) Bianchi identities.

For Noether herself, the impetus for the paper arose from the discussions over the status of energy conservation in generally covariant theories between Hilbert, Klein and Einstein, during which Hilbert commented that energy conservation for the matter fields no longer has the same status in generally covariant theories as it had in previous (non-generally covariant) theories, because it follows independently of the field equations for the matter fields. Noether's two theorems can be used to support this conjecture (see [Brading, 2005]). The discussion over the status of energy conservation in General Relativity continues, the root of the issue being that energy-momentum cannot, in general, be defined locally.⁶⁶

Today, the significance of Noether's results lie in their generality. Many of the specific connections between global spacetime symmetries and their associated conserved quantities were known before Noether's 1918 paper, and both Einstein and Hilbert anticipated some aspects of the second theorem in their investigations of energy conservation during and after the development of GTR.⁶⁷ However, her systematic treatment allows us to understand that these relations do not rely on the detailed dynamics of a particular theory, but in fact follow from the structure of Lagrangian theories and significantly weaker stipulations than the full dynamics of the theory. For example, general covariance leads to energy conservation in GTR given satisfaction of the gravitational field equations, but *independently*

⁶⁴Whether the dependencies expressed by the second theorem are trivial or not depends on the status of the fields with respect to which the local symmetry holds. It is in this way that Noether's second theorem can be used as a tool in the attempt to demarcate 'true gauge theories' from theories where the local symmetry is a 'mere mathematical artefact' (see Section 6.3 above, and [Earman, 2006]). For a 'true gauge theory' the dependencies have significant physical implications.

⁶⁵[Hilbert, 1915].

⁶⁶The energy-momentum conservation law in General Relativity is formulated in terms of the vanishing of the *covariant* divergence of the energy-momentum tensor associated with the matter fields. Alternatively, we can express this in terms of the vanishing of the coordinate divergence of the energy-momentum of the matter fields plus that of the gravitational field. The latter term falling under the divergence operator is not uniquely defined and, pertinent the issue of non-localizability, may vanish in some coordinate systems and not in others. We can understand this coordinate dependence by reflecting on the equivalence principle, according to which partitioning the inertial-gravitational field to obtain a division between inertial and gravitational forces is itself a coordinate-dependent issue. For further discussion see, for example, [Misner, Thorne and Wheeler, 1970, 467–8], and [Wald, 1984, 70]. See also [Malament, this vol., ch. 3].

⁶⁷On Einstein, see [Janssen, 2005, 75–82]; and see [Sauer, 1999] on Hilbert.

of the detailed form of those equations, and *independently* of the field equations for the matter fields (indeed, independently of whether the matter fields satisfy Euler-Lagrange equations at all).⁶⁸ Noether’s theorems are a powerful tool for investigating the structure of theories — which assumptions are required to generate which aspects of the theory, and so forth.⁶⁹

It is worthwhile mentioning a third theorem, connected with Noether’s two theorems and derived in the same context (i.e. the study of generally covariant theories of gravitation and conservation of energy) by Felix Klein [1918]. We call it the ‘Boundary theorem’ for reasons associated with its method of derivation.⁷⁰ As with Noether’s second theorem, the Boundary theorem concerns local symmetries, and results in a series of identities (termed the ‘cascade equations’ by Julia and Silva [1998]).⁷¹ We state here a simplified version of the Boundary theorem in which the action is left unchanged by an infinitesimal symmetry transformation (i.e. we do not allow for the possibility of a surface term).⁷²

The Boundary theorem (restricted form)

If a continuous group of transformations depending smoothly on ρ arbitrary functions of time and space $p_k(x)$ ($k = 1, 2, \dots, \rho$) and their first derivatives is a Noether symmetry group⁷³ of the Euler-Lagrange equations associated with a Lagrangian $L(\phi_i, \partial_\mu \phi_i, x^\mu)$, then the following three sets of ρ relations are satisfied, one for every parameter on which the symmetry group depends:

$$\sum_i \partial_\mu (b_{ki}^\mu E_i^L) = \partial_\mu j_k^\mu \tag{5}$$

$$\sum_i (b_{ki}^\mu E_i^L) = j_k^\mu - \sum_i \left[\partial_\nu \left(\frac{\partial L}{\partial(\partial_\nu \phi_i)} b_{ki}^\mu \right) \right] \tag{6}$$

$$\left(\frac{\partial L}{\partial(\partial_\mu \phi_i)} b_{ki}^\nu \right) + \left(\frac{\partial L}{\partial(\partial_\nu \phi_i)} b_{ki}^\mu \right) = 0. \tag{7}$$

Once again, the b_{ki}^ν depend on the particular transformations of the fields in question, the details of which need not concern us here. The first identity is

⁶⁸See [Brading and Brown, 2007].

⁶⁹For a discussion of this in the case of general covariance, see [Brading and Brown, 2002].

⁷⁰The Boundary theorem also appears in the work of Hermann Weyl, specialized to the case of his unified field theory (see [Weyl, 1922, 287–289]; the first appearance was in the 1919 third edition), and was published in a non-theory-specific form by Utiyama [1956; 1959].

⁷¹As with Noether’s second theorem, the Boundary theorem is a useful tool in the attempt to demarcate ‘true gauge theories’ from theories where the local symmetry is a ‘mere mathematical artefact’, through inspection of the identities that result from the theorem, and through the physical significance — or otherwise — of these identities.

⁷²For further details of the Boundary theorem, including the generalization that allows for a surface term, see [Brading and Brown, 2007].

⁷³The Boundary theorem is here stated in a restricted form such that the Noether symmetry group must belong to the restricted class of such groups associated with an *invariant* action.

connected to the existence of *superpotentials* associated with local symmetries.⁷⁴ The second equation can be used to investigate the relationship between a field and its sources. For example, in the case of classical electromagnetism, we can investigate the relationship between the local gauge symmetry of the theory and the condition that:

$$j^\mu = \partial_\nu F^{\mu\nu}, \quad (8)$$

i.e. that Maxwell's equations with dynamical sources hold. Using the case of classical electromagnetism as our example once again, the third equation becomes the condition that the electromagnetic tensor be antisymmetric (showing the relationship between this condition and the local gauge symmetry of that theory):

$$F^{\mu\nu} + F^{\nu\mu} = 0. \quad (9)$$

These remarks have been necessarily brief, and the reader is referred to [Barbashov and Nesterenko, 1983], along with Brading and Brown [2003; 2007], for detailed derivations and discussion of these results. The identities of the Boundary theorem and of Noether's two theorems are not all independent of one another, and which is most useful depends on the context and the question under consideration. As with Noether's theorems, the Boundary theorem holds independently of the specific details of the dynamical equations, and together they allow us to investigate structural features of our theories that are associated with the symmetry properties of those theories.

8 THE INTERPRETATION OF SYMMETRIES IN CLASSICAL PHYSICS

In what follows, we begin with 'Wigner's hierarchy', which has become the canonical view of the relationship between symmetries, laws and events. We supplement this with a brief discussion of the connection between symmetry and irrelevance, and how this bears on the interpretation of the various symmetries described in Section 4.2, above.

The general interpretation of symmetries in physical theories can adopt a number of complementary approaches. We can ask about the different *roles* that various symmetries play; about the epistemological, ontological or other *status* that various symmetries have; and about the significance of the structures left invariant by symmetry transformations. We end with some remarks on each of these issues.

8.1 Wigner's hierarchy

The starting point for contemporary philosophical discussion of the status and significance of symmetries in physics is Wigner's 1949 paper 'Invariance in Physical Theory', along with his three later papers published in 1964.⁷⁵ In these papers,

⁷⁴See, for example, [Trautman, 1962, 179].

⁷⁵Wigner's papers can be found in the collection *Symmetries and Reflections* [Wigner, 1967].

Wigner makes the distinction mentioned above (see Section 4.2) between geometrical and dynamical symmetries, which we will return to below. He also presents his view of the *hierarchy* of physical knowledge, according to which symmetries are viewed as properties of laws:

There is a strange hierarchy in our knowledge of the world around us. Every moment brings surprises and unforeseeable events — truly the future is uncertain. There is, nevertheless, a structure in the events around us, that is, correlations between the events of which we take cognizance. It is this structure, these correlations, which science wishes to discover, or at least the precise and sharply defined correlations. . . . We know many laws of nature and we hope and expect to discover more. Nobody can foresee the next such law that will be discovered. Nevertheless, there is a structure in the laws of nature which we call the laws of invariance. This structure is so far-reaching in some cases that laws of nature were guessed on the basis of the postulate that they fit into the invariance structure. . . . This then, the progression from events to laws of nature, and from laws of nature to symmetry or invariance principles, is what I meant by the hierarchy of our knowledge of the world around us. [Wigner, 1967, 28–30].

This view of symmetries, as properties of laws, has become canonical.

8.2 *Symmetry and irrelevance*

There is a general property of laws, or of the underlying events, to which symmetries are connected: the *irrelevance* of certain quantities that might otherwise be thought to have physical significance.⁷⁶ In Section 4.2 we outlined the variety of symmetries found in physics, and in each case the symmetry is associated with a property that is deemed irrelevant for the purposes of describing the law-governed behaviour of a system. For example, left-right symmetry means that whether a system is left-handed or right-handed is irrelevant to its law-governed evolution. Famously, this symmetry is violated in the weak interaction: the law-governed behaviour of systems turns out to be sensitive to handedness for certain processes (see [Pooley, 2003]).

In Section 4.2 we characterized the distinction between global and local symmetries mathematically, in terms of the dependence on constant parameters and arbitrary functions of time (and space) respectively. The physical meaning of this distinction can be understood through the associated properties that are deemed *irrelevant*. A *global* symmetry reflects the irrelevance of absolute values of a certain quantity: only relative values are relevant. So in Newtonian mechanics, for example, spatial translation invariance holds and absolute position is irrelevant to

⁷⁶For an analysis of the connection between symmetry, equivalence and irrelevance, see [Castellani, 2003].

the behaviour of systems.⁷⁷ Only relative positions matter, and this is reflected in the structure of the theory through the equations being invariant under global spatial translations — the equations do not depend upon, or invoke, a background structure of absolute positions.

A global symmetry is a special case of a local symmetry. A *local* symmetry reflects the irrelevance not only of absolute values, but furthermore of relative values specified at-a-distance: only *local* relative values (i.e. relative values specified at a point) are relevant. This is reflected in the structure of the theory by the equations of motion not depending upon some background structure that determines relative values at-a-distance (i.e. there is no global background structure associated with the property in question).⁷⁸

8.3 Roles of symmetries

The various different roles in which symmetries are invoked in physics have become much more evident with the advent of quantum theory.⁷⁹ Nevertheless, already with the classification of crystals using their remarkable and varied symmetry properties, we see the powerful *classificatory* role at work. Indeed, it was with René-Just Haüy's use of symmetries in this way that crystallography emerged in 1801 as a discipline distinct from mineralogy.⁸⁰ Furthermore, the *heuristic* and/or *normative* role is clear for the principle of relativity in the construction of both Special and General Relativity (see above, Section 5). The *unificatory* role, so prominent now in the attempts to unify the fundamental forces, was already present (although differing methodologically somewhat) in Hilbert's attempt to construct a generally covariant theory of gravitation and electromagnetism (see [Sauer, 1999]) and in Weyl's 1918 unified theory of gravitation and electromagnetism, for example. Symmetries may also be invoked in a variety of *explanatory* roles. For example, on the basis of Noether's first theorem (see Section 7) we might say that it is *because* of the translational symmetry of classical mechanics (plus

⁷⁷We are considering here Newtonian mechanics, without Newton's absolute space.

⁷⁸Instead, we require the explicit appearance of a *connection* in our theory, which provides the rules by which two distant objects may be brought together so that comparisons between them may be made locally.

⁷⁹The application of the theory of groups and their representations for the exploitation of symmetries in the quantum mechanics of the 1920s represents a dramatic step-change in the significance of symmetries in physics, with respect to both the foundations and the phenomenological interpretation of the theory. As Wigner emphasized on many occasions, one essential reason for the 'increased effectiveness of invariance principles in quantum theory' [Wigner, 1967, 47] is the linear nature of the state space of a quantum physical system, corresponding to the possibility of superposing quantum states. For details on the application of symmetries in quantum physics we refer the reader to [Dickson, this vol., ch. 4, Section 3.3], [Landsman, this vol., ch. 5, Section 4.1], and [Halvorson, this vol., ch. 8, Section 5.2]. For philosophical discussions see [Brading and Castellani, 2003].

⁸⁰The use of discrete symmetries in crystallography continued through the nineteenth century in the work of J. F. Hessel and A. Bravais, leading to the 32 point transformation crystal classes and the 14 Bravais lattices. These were combined into the 230 space groups in the 1890s by E. S. Fedorov, A. Schönflies, and W. Barlow. The theory of discrete groups continues to be important in such fields as solid state physics, chemistry, and materials science.

satisfaction of other conditions) that linear momentum is conserved in that theory. Another example would be an appeal to symmetry principles as an explanation, via Wigner’s hierarchy, for (i) aspects of the *form* of the laws, and thereby (ii) why certain events occur and others do not.

8.4 *Status of symmetries*

Are symmetries ontological, epistemological, or methodological in status? It is clear that symmetries have an important heuristic function, as discussed above (Section 5) in the context of relativity. This indicates a methodological status, something that becomes further developed within the context of quantum theory. We can also ask whether we should attribute an ontological or epistemological status to symmetries.

According to an ontological viewpoint, symmetries are seen as “existing in nature”, or characterizing the structure of the physical world. One reason for attributing symmetries to nature is the so-called geometrical interpretation of spatiotemporal symmetries, according to which the spatiotemporal symmetries of *physical laws* are interpreted as symmetries of *spacetime itself*, the “geometrical structure” of the physical world. Moreover, this way of seeing symmetries can be extended to non-external symmetries, by considering them as properties of other kinds of spaces, usually known as “internal spaces”.⁸¹ The question of exactly what a realist would be committed to on such a view of internal spaces remains open, and an interesting topic for discussion.

One approach to investigating the limits of an ontological stance with respect to symmetries would be to investigate their empirical or observational status: can the symmetries in question be directly observed? We first have to address what it means for a symmetry to be observable, and indeed whether all symmetries have the same observational status. Kosso [2000] arrives at the conclusion that there are important differences in the empirical status of the different kinds of symmetries. In particular, while global continuous symmetries can be directly observed — via such experiments as the Galilean ship experiment — a local continuous symmetry can have only indirect empirical evidence.⁸²

The direct observational status of the familiar global spacetime symmetries leads us to an epistemological aspect of symmetries. According to Wigner, the spatiotemporal invariance principles play the role of a prerequisite for the very possibility of discovering the laws of nature: ‘if the correlations between events changed from day to day, and would be different for different points of space, it would be impossible to discover them’ [Wigner, 1967]. For Wigner, this conception of symmetry principles is essentially related to our ignorance (if we could directly know all the laws of nature, we would not need to use symmetry principles in our search for them). Such a view might be given a methodological interpretation, ac-

⁸¹See Section 4.2, above, for the varieties of symmetry.

⁸²See Section 6.1, above; and Brading and Brown [2003b], who argue for a different interpretation of Kosso’s examples.

according to which such spatiotemporal regularities are presupposed in order for the enterprise of discovering the laws of physics to get off the ground.⁸³ Others have arrived at a view according to which symmetry principles function as “transcendental principles” in the Kantian sense (see for instance [Mainzer, 1996]). It should be noted in this regard that Wigner’s starting point, as quoted above, does not imply exact symmetries — all that is needed epistemologically (or methodologically) is that the global symmetries hold approximately, for suitable spatiotemporal regions, so that there is sufficient stability and regularity in the events for the laws of nature to be discovered.

As this discussion, and that of the preceding Subsections, indicate, the *differences* between various types of symmetry become important before we have ventured very far into interpretational issues. For this reason, much recent work on the interpretation of symmetry in physical theory has focussed not on general questions, such as those sketched above, but on addressing interpretational questions specific to particular symmetries.⁸⁴

8.5 *Symmetries, objectivity, and objects*

Turning now to the issue of the structures left invariant by symmetry transformations, the old and natural idea that what is objective should not depend upon the particular perspective under which it is taken into consideration is reformulated in the following group theoretical terms: what is objective is what is invariant with respect to the relevant transformation group. This connection between symmetries and objectivity is something that has a long history going back to the early twentieth century at least. It was highlighted by Weyl [1952], where he writes that ‘We found that objectivity means invariance with respect to the group of automorphisms.’ This connection between objectivity and invariance was discussed particularly in the context of Relativity Theory, both Special and General. We recall Minkowski’s famous phrase ([1908] 1923, 75) that ‘Henceforth space by itself, and time by itself, are doomed to fade away into mere shadows, and only a kind of union of the two will preserve an independent reality’, following his geometrization of Einstein’s Special Theory of Relativity, and the recognition of the spacetime interval (rather than intervals of space and of time) as the geometrically invariant quantity. The connection between objectivity and invariance in General Relativity was discussed by, amongst others, Hilbert and Weyl, and continues to be an issue today.⁸⁵

⁸³We are grateful to Brandon Fogel for this point, and for the comparison he suggested between this view of spatiotemporal symmetries and the methodological face of Einstein’s notion of separability.

⁸⁴These include the varieties of gauge invariance found in classical electromagnetism and in quantum theories, along with general covariance in GTR (these being continuous symmetries), plus the discrete symmetries of parity (violated in the weak interaction) and permutation invariance, both of which are found in classical theory but require reconsideration in the light of quantum theory. See [Brading and Castellani, 2003].

⁸⁵We saw above (Sections 6.2 and 6.3) some aspects of this debate in the discussion of Einstein’s ‘hole argument’ and of the status of observables in GTR.

Related to this is the use of symmetries to characterize the *objects* of physics as sets of invariants. Originally developed in the context of quantum theory, this approach can also be applied in classical physics.⁸⁶ The basic idea is that the invariant quantities — such as mass and charge — are those by which we characterize objects. Thus, through the application of group theory we can use symmetry considerations to determine the invariant quantities and “construct” or “constitute” objects as sets of these invariants.⁸⁷

In conclusion, then, the philosophical questions associated with symmetries in classical physics are wide-ranging. What we have offered here is nothing more than an overview, influenced by our own interests and puzzles, which we hope will be of service in further explorations of this philosophically and physically rich field.

ACKNOWLEDGEMENTS

We are grateful to the editors, Jeremy Butterfield and John Earman, for their encouragement and detailed comments. We would also like to thank Brandon Fogel, Brian Pitts, and Thomas A. Ryckman for their comments and suggestions.

BIBLIOGRAPHY

- [Adler *et al.*, 1975] R. Adler, M. Bazin, and M. Schiffer. *Introduction to General Relativity*. 2nd edition, McGraw-Hill Book Company, New York, St Louis, San Francisco, Toronto, London, Sydney, 1975.
- [Anderson, 1967] J. L. Anderson. *Principles of Relativity Physics*. New York and London: Academic press, 1967.
- [Barbashov and Nesterenko, 1983] B. M. Barbashov and V. V. Nesterenko. Continuous symmetries in field theory, *Fortschr. Phys.*, 31: 535-67, 1983.
- [Barbour, 1989] J. B. Barbour. *Absolute or Relative Motion?* Vol. 1 *The Discovery of Dynamics*. Oxford University Press, 1989.
- [Barbour, 2005] J. B. Barbour. *Absolute or Relative Motion?* Vol. 2 *The Deep Structure of General Relativity*. Oxford University Press, 2005.
- [Belot, 1998] G. Belot. Understanding electromagnetism. *British Journal for the Philosophy of Science*, 49: 531–55, 1998.
- [Belot, 2006] G. Belot. The representation of time and change in mechanics. This volume.
- [Brading, 2005] K. Brading. A note on general relativity, energy conservation, and Noether’s theorems. In J. Eisenstaedt and A. J. Kox, Birkhäuser (eds.) *The Universe of General Relativity (Einstein Studies)*, 2005.
- [Brading and Brown, 2002] K. Brading and H. R. Brown. General covariance from the perspective of Noether’s theorems. *Diálogos*: 59-86, 2002.
- [Brading and Brown, 2003] K. Brading and H. R. Brown. Symmetries and Noether’s theorems. In K. Brading and E. Castellani (eds.), pages 89-109, 2003.
- [Brading and Brown, 2004] K. Brading and H. R. Brown. Are gauge symmetry transformations observable?. *British Journal for the Philosophy of Science*, 55: 645-65, 2004.
- [Brading and Brown, 2007] K. Brading and H. R. Brown. Noether’s theorems, gauge symmetries and general relativity. Manuscript, 2007.
- [Brading and Castellani, 2003] K. Brading and E. Castellani (eds.). *Symmetry in Physics: Philosophical Reflections*. Cambridge University Press, 2003.

⁸⁶See Max Born, reprinted in [Castellani, 1998].

⁸⁷For further discussion see [Castellani, 1998, part II].

- [Brading and Castellani, 2006] K. Brading and E. Castellani. Curie's Principle, Encore. In preparation, 2006.
- [Brading and Ryckman, 2007] K. Brading and T. A. Ryckman. Hilbert's axiomatic method and the foundations of physics: an interpretation of generally covariant physics and a revision of Kant's epistemology. Manuscript, 2007.
- [Brown, 2006] H. Brown. *Physical Relativity: Spacetime Structure from a Dynamical Perspective*. Oxford University Press, 2006.
- [Butterfield, 2004] J. Butterfield. Between laws and models: Some philosophical morals of Lagrangian mechanics; available at Los Alamos arXive: <http://arxiv.org/abs/physics/0409030>; and at Pittsburgh archive: <http://philsci-archive.pitt.edu/archive/00001937/>.
- [Butterfield, 2006] J. Butterfield. On symmetry and conserved quantities in classical mechanics. Forthcoming in W. Demopoulos and I. Pitowsky (eds.), *Festschrift for Jeffrey Bub*. Kluwer: University of Western Ontario Series in Philosophy of Science, 2006; available at Los Alamos arXive: <http://arxiv.org/abs/physics/>; and at Pittsburgh archive: <http://philsci-archive.pitt.edu/archive/00002362/>.
- [Butterfield, 2006] J. Butterfield. On symplectic reduction in classical mechanics. This volume.
- [Castellani, 1998] E. Castellani (ed.). *Interpreting Bodies. Classical and Quantum Objects in Modern Physics*. Princeton University Press, 1998.
- [Castellani, 2003] E. Castellani. Symmetry and equivalence. In K. Brading and E. Castellani (eds.), pages 321-334, 2003.
- [Chalmers, 1970] A. F. Chalmers. Curie's principle. *The British Journal for the Philosophy of Science*, 21: 133-148, 1970.
- [Corry, 2004] L. Corry. *David Hilbert and the Axiomatization of Physics (1898-1918)*. Dordrecht: Kluwer Academic, 2004.
- [Curie, 1894] P. Curie. Sur la symétrie dans les phénomènes physiques. Symétrie d'un champ électrique et d'un champ magnétique. *Journal de Physique*, 3^e série, vol. 3: 393-417, 1894.
- [Curie, 1981] P. Curie. On symmetry in physical phenomena. Trans. J. Rosen and P. Copié, *Am. J. Phys.*, 49(4): 17-25, 1981.
- [Descartes, 1998] R. Descartes. *The World and other Writings*. S. Gaukroger (ed. and trans.), Cambridge University Press, 1998.
- [Dickson, 2006] M. Dickson. Non-relativistic quantum theory. This volume, 2006.
- [Doughty, 1990] N. Doughty. *Lagrangian Interaction*. Addison-Wesley, 1990.
- [Earman, 2004] J. Earman. Curie's principle and spontaneous symmetry breaking. *International Studies in Philosophy of Science*, 18: 173-199, 2004.
- [Earman, 2006a] J. Earman. Two challenges to the substantive requirement of general covariance. *Synthese*, in press, 2006.
- [Earman, 2006b] J. Earman. Aspects of determinism in modern physics. This volume.
- [Halvorson, 2006] H. Halvorson. Algebraic quantum field theory. This volume.
- [Hawkins, 2000] T. Hawkins. *Emergence of the Theory of Lie Groups: an Essay in the History of Mathematics 1869-1926*, New York: Springer, 2000.
- [Hilbert, 1915] D. Hilbert. Die Grundlagen der Physik (Erste Mitteilung), Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen. *Mathematisch-physikalische Klasse*: 395-407, 1915.
- [’t Hooft, 2006] ’t Hooft. The conceptual basis of quantum field theory. This volume.
- [Howard, 2007] D. Howard. 'And I shall not mingle conjectures and certainties': The roots of Einstein's principle theories/constructive theories distinction, manuscript, 2007.
- [Ismael, 1997] J. Ismael. Curie's principle. *Synthese*, 110: 167-190, 1997.
- [Janssen, 2005] M. Janssen. Of pots and holes: Einstein's bumpy road to general relativity. *Ann. Phys.*, 14, Supplement: 58-85, 2005.
- [Julia and Silva, 1998] B. Julia and S. Silva. Current and superpotentials in classical gauge invariant theories I. Local results with applications to perfect fluids and general relativity, 1998. gr-qc/9804029 v2.
- [Klein, 1918] F. Klein. Über die Differentialgesetze für die Erhaltung von Impuls und Energie in der Einsteinschen Gravitationstheorie. Königliche Gesellschaft der Wissenschaften zu Göttingen. *Mathematisch-physikalische Klasse. Nachrichten*: 469-92, 1918.
- [Kosso, 2000] P. Kosso. The empirical status of symmetries in physics. *British Journal for the Philosophy of Science*, 51: 81-98, 2000.

- [Lanczos, 1949] C. Lanczos. *The Variational Principles of Mechanics*. University of Toronto Press, [1949], 1962.
- [Landsman, 2006] N. P. Landsman. Between classical and quantum. This volume.
- [Mainzer, 1996] K. Mainzer. *Symmetries of Nature. A Handbook*. New York: de Gruyter, 1996.
- [Malament, 2006] D. Malament. Classical relativity theory. This volume.
- [Miller, 1981] A. I. Miller. *Albert Einstein's Special Theory of Relativity*. Addison-Wesley Publishing Company, 1981.
- [Minkowski, 1908] H. Minkowski. Space and time. In *The Principle of Relativity. A Collection of Original Memoirs on the Special and General Theory of Relativity*. New York: Dover, pages 75-91, [1908], 1923.
- [Misner *et al.*, 1970] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. New York: W. H. Freeman and Company, 1970.
- [Noether, 1971] E. Noether. Noether's theorem. *Transport Theory and Statistical Physics*, 1: 183-207, 1971.
- [Norton, 1984] J. Norton. How Einstein found his field equations: 1912-1915. *Hist. Stud. Phys. Sci.*, 14: 253-316, 1984.
- [Norton, 1993] J. Norton. General covariance and the foundations of general relativity: eight decades of dispute. *Rep. Prog. Phys.*, 56: 791-858, 1993.
- [Norton, 2003] J. Norton. General covariance, gauge theories, and the Kretschmann objection. In K. Brading and E. Castellani (eds.), pages 110-123, 2003.
- [Ohanian and Ruffini, 1994] H. C. Ohanian and R. Ruffini. *Gravitation and Spacetime*, 2nd edition. London, New York: W. W. Norton and company, 1994.
- [Olver, 1993] P. J. Olver. *Applications of Lie Groups to Differential Equations*, 2nd edition, New York: Springer, 1993.
- [O'Raifeartaigh, 1997] L. O'Raifeartaigh. *The Dawning of Gauge Theory*. Princeton University Press, 1997.
- [Pitts, 2006] J. B. Pitts. Absolute objects and counterexamples: Jones-Geroch Dust, Torretti constant curvature, Tetrad-Spinor, and scalar density. *Studies in History and Philosophy of Modern Physics*, forthcoming, 2006. Manuscript available at Los Alamos arXiv: <http://arxiv.org/abs/gr-qc/0506102/>.
- [Pooley, 2003] O. Pooley. Handedness, parity violation, and the reality of space. In K. Brading and E. Castellani (eds.), pages 250-280, 2003.
- [Radicati, 1987] L. A. Radicati. Remarks on the early developments of the notion of symmetry breaking. In M. G. Doncel, A. Hermann, L. Michel, and A. Pais (eds.), *Symmetries in physics (1600-1980)*. Barcelona: Servei de Publicaciones, pages 195-206, 1987.
- [Rovelli, 2006] C. Rovelli. Quantum gravity. This volume.
- [Ryckman, 2005] T. Ryckman. *The Reign of Relativity*. Oxford University Press, 2005.
- [Sauer, 1999] T. Sauer. The relativity of discovery: Hilbert's first note on the foundations of physics. *Arch. Hist. Exact Sci.*, 53: 529-575.
- [Shubnikow and Koptsik, 1974] A. V. Shubnikow and V. A. Koptsik. *Symmetry in Science and Art*. London: Plenum Press, 1974.
- [Stachel, 1993] J. Stachel. The meaning of general covariance: the hole story. In J. Earman *et al.* (eds.), *Philosophical Problems of the Internal and External World: Essays on the Philosophy of Adolf Grünbaum..* University of Pittsburgh Press, pages 129-160, 1993.
- [Torretti, 1983] R. Torretti. *Relativity and Geometry*. Oxford: Pergamon Press, 1983.
- [Trautman, 1962] A. J. Trautman. Conservation laws in general relativity. In L. Witten (ed.), *Gravitation: An Introduction to Current Research*. New York: John Wiley and Sons, pages 169-98, 1962.
- [Utiyama, 1956] R. Utiyama. Invariant theoretical interpretation of interaction. *Physical Review*, 101: 1597-607, 1956.
- [Utiyama, 1959] R. Utiyama. Theory of invariant variation in the generalized canonical dynamics. *Prog. Theor. Phys. Suppl.*, 9: 19-44, 1959.
- [Wald, 1984] R. Wald. *General Relativity*. University of Chicago Press, 1984.
- [Weyl, 1922] H. Weyl. *Space, Time, Matter*. Reprinted by Dover (1952), 1922.
- [Weyl, 1952] H. Weyl. *Symmetry*. Princeton University Press, 1952.
- [Whittaker, 1904] E. T. Whittaker. *A Treatise on the Analytical Dynamics of Particles and Rigid Bodies*. Cambridge University Press, ([1904] 1989).
- [Wigner, 1967] E. P. Wigner. *Symmetries and Reflections*. Bloomington: Indiana University Press, 1967.

[Yaglom, 1988] I. M. Yaglom. *Felix Klein and Sophus Lie. Evolution of the Idea of Symmetry in the Nineteenth Century*. Boston-Basel: Birkhäuser, 1988.

This page intentionally left blank

ASPECTS OF DETERMINISM IN MODERN PHYSICS

John Earman

1 INTRODUCTION

The aims of this chapter are to review some aspects of determinism that are familiar to physicists but are little discussed in the philosophical literature and to show how these aspects connect determinism to issues about symmetries in physics, the structure and ontological status of spacetime, predictability, and computability.¹ It will emerge that in some respects determinism is a robust doctrine and is quite hard to kill, while in other respects it is fragile and requires various enabling assumptions to give it a fighting chance. It will also be seen that determinism is far from a dead issue. Whether or not ordinary non-relativistic quantum mechanics (QM) admits a viable deterministic underpinning is still a matter of debate. Less well known is the fact that in some cases QM turns out to be more deterministic than its classical counterpart. Quantum field theory (QFT) assumes determinism, at least at the classical level, in order to construct the field algebra of quantum observables. Determinism is at the heart of the cosmic censorship hypothesis, the most important unsolved issue in classical general relativity theory (GTR). And issues about the nature and status of determinism lie at the heart of key foundation issues in the search for a theory of quantum gravity.

2 PRELIMINARIES

2.1 The metaphysics of determinism

The proposal is to begin by getting a grip on the doctrine of determinism as it was understood pre-GTR and pre-QM, and then subsequently to try to understand how the doctrine has to be adjusted to accommodate these theories. In pre-GTR physics, spacetime serves as a fixed background against which the drama of physics is enacted. In pre-QM physics it was also assumed that there is a set \mathcal{O} of genuine physical magnitudes (a.k.a. “observables”) each of which takes a determinate

¹Recent surveys of determinism are found in Butterfield [1998], Earman [2004a], and Hoefer [2003]. A collection of articles on various aspects of determinism is found in Atmanspacher and Bishop [2002].

value at every moment of time; call these the *occurrent* magnitudes. Other physical magnitudes may be dispositional in character and may take on determinate values only in appropriate contexts; but it was assumed that these dispositional magnitudes supervene on the nondispositional magnitudes.² A *history* H is a map from \mathbb{R} to tuples of values of the basic magnitudes, where for any $t \in \mathbb{R}$ the *state* $H(t)$ gives a snapshot of behavior of the basic magnitudes at time t . The world is *Laplacian deterministic* with respect to \mathcal{O} just in case for any pair of histories H_1, H_2 satisfying the laws of physics, if $H_1(t) = H_2(t)$ for some t , then $H_1(t) = H_2(t)$ for all t .

Several remarks are in order. First, the ‘ t ’ which appears in the above definition is supposed to be a *global time function*. This notion can be defined in a manner that applies to classical, special relativistic, and general relativistic spacetimes: a global time function is a smooth map $t : \mathcal{M} \rightarrow \mathbb{R}$, where \mathcal{M} is the spacetime manifold, such that for any $p, q \in \mathcal{M}$, $t(p) < t(q)$ just in case there is a future directed timelike curve from p to q .³ In classical spacetimes, all of which possess an absolute (or observer independent) notion of simultaneity, a timelike curve is one which is oblique to the planes of absolute simultaneity. And the levels $t = \text{const}$ of a global time function must coincide with the planes of simultaneity; thus, in the classical setting t is determined up to a transformation of the form $t \rightarrow t' = t'(t)$. In the relativistic setting a timelike curve is one whose tangent at any point lies inside the light cone at that point. In causally pathological general relativistic spacetimes (e.g. Gödel spacetime — see Section 6.1) there can be no global time function, and the global sense of Laplacian determinism as defined above makes no sense.⁴ But if one global time function exists for a relativistic spacetime, then many exist. A poor choice of global time function can lead to the failure of Laplacian determinism on the above definition. Thus, in the relativistic setting, the definition of determinism must be construed as applying to a suitable choice of time function, the nature of which will be clarified below.

Second, the above formulation of determinism assumes a distinction between laws of nature on one hand and initial/boundary conditions on the other. Where this distinction becomes mushy, so does the doctrine of determinism. There is a

²The general idea of supervenience is that X supervenes on Y iff within the range of possible cases, there is no difference in X without a difference in Y . The strength and type of supervenience depends on what are counted as possible cases. Here the concern is mainly with physical supervenience where the possible cases are those compatible with the laws of physics.

³This definition presupposes that the spacetime is temporally orientable and that one of the orientations has been singled out as giving the future direction of time. The first presupposition is satisfied for classical and special relativistic spacetimes. A general relativistic spacetime (see [Malament, this vol.]) may not be temporally orientable, but a covering spacetime always is since temporal orientability fails only if the spacetime is not simply connected. The second presupposition implies that some solution to the problem of the direction of time has been found (see [Uffink, this vol.]).

⁴A necessary and sufficient condition for the existence of a global time function for a relativistic spacetime is *stable causality* which (roughly speaking) requires that there exists a widening of the null cones that does not result in closed timelike curves; for a precise definition, see [Wald, 1984, 198-199]. Not only does Gödel spacetime not admit a global time function, it does not admit any global time slices (i.e. spacelike hypersurfaces without edges).

huge philosophical literature on laws of nature.⁵ Since most of it is unilluminating when it comes to understanding the nature and function of laws in the practice of physics, it will be largely ignored here. For present purposes I will simply stipulate that an acceptable account of laws must satisfy the empiricist constraint that the laws supervene on the totality of non-modal, particular facts.⁶ Philosophers like to speculate about non-empiricist laws; but such entities, should they exist, would seem to be beyond the ken of science, and as such they are irrelevant for present purposes. I prefer David Lewis' [1973, 72-77] way of fulfilling the empiricist constraint since it connects the account of laws to the practice of physics: the laws of physics are the axioms or postulates that appear in the ideal theory of physics, where the ideal theory is the one that, within the class of true theories, achieves the best balance between simplicity and information content. All of the historical examples we have of candidates for deterministic laws involve a relatively small subset $\mathcal{B} \subset \mathcal{O}$ of basic occurrent magnitudes, the assumption being that the remaining ones supervene on those of \mathcal{B} .⁷ This is hardly surprising if, as has been claimed, simplicity is a crucial feature of physical laws. Hermann Weyl shared the conviction that simplicity must figure into an account of laws, but he noted that "this circumstance is apt to weaken the metaphysical power of determinism, since it makes the meaning of natural law depend on the fluctuating distinction between simple and complicated functions or classes of functions" [1932, 42]. This is, I think, a consequence that has to be swallowed and digested. Philosophers who are answerable only to their armchairs are free to think otherwise.

Third, it is conceptually possible that the world could be partially deterministic, i.e. deterministic with respect to partial histories defined by the values of magnitudes in some proper subset $\mathcal{D} \subset \mathcal{O}$ of the occurrent physical magnitudes but non-deterministic with respect to partial histories defined by the values of magnitudes in some other proper subset $\mathcal{N} \subset \mathcal{O}$. But it is hard to imagine a scenario in which this could happen if both \mathcal{D} and \mathcal{N} are basic magnitudes. For in order that the non-deterministic evolution of the elements \mathcal{N} not upset deterministic evolution for \mathcal{D} , the magnitudes in \mathcal{N} must not interact with those in \mathcal{D} , or else there would have to be a conspiracy in which the upsetting effects of the \mathcal{N} magnitudes on \mathcal{D} cancel out, which is operationally the same. However, this plausibility consideration fails to operate when the \mathcal{N} are non-basic magnitudes; in particular, as discussed below, stochastic processes on one level can supervene on deterministic processes at a lower level (see [Butterfield, 1998]). This fact makes the inference from observed stochastic behavior to indeterminism fraught with peril.

Fourth, the laws of physics typically take the form of differential equations, in which case the issue of Laplacian determinism translates into the question of whether the equations admit an *initial value formulation*, i.e. whether for ar-

⁵For an overview of different accounts of laws of nature, see [Carroll, 2004].

⁶This is what David Lewis has termed "Humean supervenience" with regards to laws of nature; for a defense, see [Earman and Roberts, 2006].

⁷For example, in classical particle mechanics the elements of \mathcal{B} are the positions and momenta of the particles, and it is assumed that any other genuine mechanical magnitude can be expressed as a functional of these basic magnitudes.

bitrary initial data there exists a unique solution agreeing with the given initial data.⁸ What counts as initial data depends on the details of the case, but typically it consists of the instantaneous values of the independent variables in the equations, together with the instantaneous values of a finite number of time derivatives of these variables. “Arbitrary” initial data might be thought to include any kinematically possible values of the relevant variable — as with the initial values of particle positions and velocities in Newtonian mechanics — but “arbitrary” must be taken to mean arbitrary within the bounds of compatibility with the equations of motion, which may impose non-trivial constraints on the initial data. This leads to the next remark.

Fifth, in the relativistic setting, field equations often factor into constraint equations, which place restrictions on the initial data, and the evolution equations, which govern how initial data satisfying the constraint equations evolve over time — Maxwell’s equations for electromagnetism and Einstein’s gravitational field equations being prime examples. In these instances the evolution equations guarantee that once the constraint equations are satisfied they continue to be satisfied over time. This should be a feature of deterministic equations, for if the data at some future time in the unique solution picked out by the initial data do not satisfy the constraints, then the laws are self-undermining. It could be argued that a basic feature of time in relativistic worlds — perhaps the key feature that separates the time dimension from the space dimensions — lies precisely in this separation of evolution and constraint equations.⁹

Sixth, while there is no *a priori* guarantee that the laws of the ideal theory of physics will be deterministic, the history of physics shows that determinism is taken to be what might be termed a ‘defeasible methodological imperative’: start by assuming that determinism is true; if the candidate laws discovered so far are not deterministic, then presume that there are other laws to be discovered, or that the ones so far discovered are only approximations to the correct laws; only after long and repeated failure may we entertain the hypothesis that the failure to find deterministic laws does not represent a lack of imagination or diligence on our part but reflects the fact that Nature is non-deterministic. An expression of this sentiment can be found in the work of Max Planck, one of the founders of quantum physics: determinism (a.k.a. the law of causality), he wrote, is a “heuristic principle, a signpost and in my opinion the most valuable signpost we possess, to guide us through the motley disorder of events and to indicate the direction in which scientific inquiry should proceed in order to attain fruitful results” [1932, 26; my translation].¹⁰

⁸And as will be discussed below, there are further issues, such as whether the solution depends continuously on the initial data.

⁹See [Callender, 2005] and [Skow, 2005] for defenses of related views on the difference between space and time.

¹⁰For a history of the debates about the status of determinism among the founding fathers of QM, see [Cushing, 1994] and [Stöltzner, 2003].

2.2 Varieties of determinism

There is a tendency in the philosophical literature to fixate on the Laplacian variety of determinism. But other kinds of determinism crop up in physics. For example, some processes are described by delay-differential equations for which instantaneous initial data may not suffice to single out a unique solution. A simple example is given by the first order ordinary differential equation (ode) $\dot{x}(t) = x(t - C)$ with a constant delay $C > 0$. Laplacian determinism fails since given initial data $x(0)$ is compatible with multiple solutions. However, a near cousin of Laplacian determinism holds since a specification of $x(t)$ for the interval of time $t \in [-C, 0]$ fixes a unique solution.¹¹ If the constant delay C replaced by a function $\tau(t)$ of t which is unbounded, or if the delay-differential equation has a more complicated form than in the simple example, then even the weakened forms of Laplacian determinism can fail. An illustration of the latter is given by the equation of motion is $\dot{x}(t) = f(t)x(t - 1)$ where $f(t)$ is a continuous function that vanishes outside of $[0, 1]$ and satisfies $\int f(t)dt = -1$. Raju [1994, 120ff] gives an example of an f such that unless $x(t)$ is identically 0 for all $t \geq 1$, the equation of motion admits no solutions for $t < 0$; whereas if $x(t)$ is identically zero for $t \geq 1$, then the equation of motion admits an infinity of solutions for $t < 0$. Changing the delay term $x(t - 1)$ in this example to an advance term $x(t + 1)$ produces an example where an entire past history fails to fix a unique future solution. Very little is known about the initial value problem for what is probably the most important physical application of delay/advance differential equations; namely, charged particles moving under their mutual retarded/advanced interactions.¹²

For sake of definiteness, fix on the Laplacian variety of determinism. Within this variety there is a distinction between future and past determinism. *Past Laplacian determinism* means that for any pair of histories H_1, H_2 satisfying the laws of physics, if $H_1(t) = H_2(t)$ for some t , then $H_1(t') = H_2(t')$ for all $t' > t$. *Future Laplacian determinism* is defined analogously. In principle, Laplacian determinism can hold in one direction of time but not in the other. However, if the laws of motion are time reversal invariant, then future and past determinism stand or fall together. Time reversal invariance is the property that if H is a history satisfying the laws, then so is the 'time reverse' history H^T , where $H^T(t) := {}^R H(-t)$ and where ' R ' is the reversal operation that is defined on a case-by-case basis, usually by analogy with classical particle mechanics where $H(t) = (\mathbf{x}(t), \mathbf{p}(t))$, with $\mathbf{x}(t)$ and $\mathbf{p}(t)$ being specifications respectively of the particle positions and momenta at t , and ${}^R H(t) = (\mathbf{x}(t), -\mathbf{p}(t))$.¹³ Since all of the plausible candidates for fundamental laws of physics, save those for the weak interactions

¹¹See [Driver, 1977] for relevant results concerning delay-differential equations.

¹²Driver [1979] studied the special case of identically charged particles confined to move symmetrically on the x -axis under half-retarded and half-advanced interactions. He showed that, provided the particles are sufficiently far apart when they come to rest, a unique solution is determined by their positions when they come to rest.

¹³A different account of time reversal invariance is given in [Albert, 2000, Ch. 1]; but see [Earman, 2002] and [Malament, 2004].

of elementary particles, are time reversal invariant, the distinction between past and future determinism is often ignored.

This is the first hint that there are interesting connections between determinism and symmetry properties.¹⁴ Many other examples will be encountered below, starting with the following section.

2.3 *Determinism and symmetries: Curie's Principle*

The statement of what is now called 'Curie's Principle' was announced in 1894 by Pierre Curie:

(CP) When certain effects show a certain asymmetry, this asymmetry must be found in the causes which gave rise to it. [Curie 1894, 401]

Some commentators see in this Principle profound truth, while others see only falsity, and still others see triviality (compare [Chalmers, 1970]; [Radicati, 1987]; [van Fraassen, 1991, 23–24], and [Ismael, 1997]). My reading of (CP) makes it a necessary truth. It takes (CP) to assert a conditional:

If

(CP1) the laws of motion governing the system are deterministic; and
 (CP2) the laws of motion governing the system are invariant under a symmetry transformation; and (CP3) the initial state of the system is invariant under said symmetry

then

(CP4) the final state of the system is also invariant under said symmetry

When the first clause (CP1) in the antecedent holds, the second clause (CP2) can be understood as follows: if an initial state is evolved for a period Δt and then the said symmetry is applied to the (unique) evolved state, the result is the same as first applying the symmetry to the initial state and evolving the resulting state for a period Δt . With this understanding, the reader can easily derive (CP4) from (CP1)-(CP3). Concrete instantiations of Curie's principle at work in classical and relativistic physics can be found in [Earman, 2004b]. An instantiation for GTR is mentioned in Section 6.3 below.¹⁵

Although (CP) is a necessary truth, it is far from a triviality since it helps to guide the search for a causal explanation of an asymmetry in what is regarded as the final state of system: either the asymmetry is already present in the initial state; or else the initial state is symmetric and the asymmetry creeps in over time,

¹⁴See [Brading and Castellani, this vol.] for a discussion of symmetries and invariances in modern physics.

¹⁵For additional remarks on Curie's principle, see [Brading and Castellani, this vol.].

either because the laws that govern the evolution of the system do not respect the symmetry or because they are non-deterministic. If, as is often the case, the latter two possibilities are ruled out, then the asymmetry in the final state must be traceable to an asymmetry in the initial state. It is also worth noting that the use of (CP) has ramifications for the never ending debate over scientific realism; for the asymmetry in the initial state may be imperceptible not only to the naked eye but to any macroscopic means of detection.¹⁶

3 DETERMINISM AND INDETERMINISM IN CLASSICAL PHYSICS

3.1 *The hard road to determinism in classical physics*

Classical physics is widely assumed to provide a friendly environment for determinism. In fact, determinism must overcome a number of obstacles in order to achieve success in this setting. First, classical spacetime structure may not be sufficiently rich to support Laplacian determinism for particle motions. Second, even if the spacetime structure is rich, uniqueness can fail in the initial value problem for Newtonian equations of motion if the force function does not satisfy suitable continuity conditions. Third, the equations of motion that typically arise for classical particles plus classical fields, or for classical fields alone, do not admit an initial value formulation unless supplementary conditions are imposed. Fourth, even in cases where local (in time) uniqueness holds for the initial value problem, solutions can break down after a finite time.

The following subsection takes up the first of these topics — the connection between determinism and the structure and ontology of classical spacetimes. The others are taken up in due course.

3.2 *Determinism, spacetime structure, and spacetime ontology*

Here is the (naive) reason for thinking that neither Laplacian determinism nor any of its cousins stands a chance unless supported by enough spacetime structure of the right kind. Assume that the (fixed) classical spacetime background is characterized by a differentiable manifold \mathcal{M} and various geometric object fields O_1, O_2, \dots, O_M on \mathcal{M} . And assume that the laws of physics take the form of equations whose variables are the O_i 's and additional object fields P_1, P_2, \dots, P_N describing the physical contents of the spacetime. (For the sake of concreteness, the reader might want to think of the case where the P_j 's are vector fields whose integral curves are supposed to be the world lines of particles.) A symmetry of the spacetime is a diffeomorphism d of \mathcal{M} onto itself which preserves the background structure given by the O_i 's — symbolically, $d^*O_i = O_i$ for all values of i , where

¹⁶For a more detailed discussion of Curie's Principle and its connection to spontaneous symmetry breaking in quantum field theory see [Earman, 2004b]; for spontaneous symmetry breaking in quantum statistical physics, see [Emch, this vol.].

d^* denotes the drag along by d .¹⁷ By the assumption on the form of the laws, a spacetime symmetry d must also be a symmetry of the laws of motion in the sense that if $\langle \mathcal{M}, O_1, O_2, \dots, O_M, P_1, P_2, \dots, P_N \rangle$ satisfies the laws of motion, then so does $\langle \mathcal{M}, O_1, O_2, \dots, O_M, d^*P_1, d^*P_2, \dots, d^*P_N \rangle$.¹⁸

Now the poorer the structure of the background spacetime, the richer the spacetime symmetries. And if the spacetime symmetry group is sufficiently rich, it will contain elements that are the identity map on the portion of spacetime on or below some time slice $t = \text{const}$ but non-identity above. We can call such a map a ‘determinism killing symmetry’ because when applied to any solution of the equations of motion, it produces another solution that is the same as the first for all past times but is different from the first at future times, which is a violation of even the weakest version of future Laplacian determinism.

As an example, take *Leibnizian spacetime*,¹⁹ whose structure consists of all and only the following: a notion of absolute or observer-independent simultaneity; a temporal metric (giving the lapse of time between non-simultaneous events); and a Euclidean spatial metric (giving the spatial distance between events lying on a given plane of absolute simultaneity). In a coordinate system (x^α, t) , $\alpha = 1, 2, 3$ adapted to this structure, the spacetime symmetries are

$$\begin{aligned} (1) \quad x^\alpha &\rightarrow x'^\alpha = R_\beta^\alpha(t)x^\beta + a^\alpha(t) & \alpha, \beta = 1, 2, 3 \\ t &\rightarrow t' = t + \text{const} \end{aligned}$$

where $R_\beta^\alpha(t)$ is an orthogonal time dependent matrix and the $a^\alpha(t)$ are arbitrary smooth functions of t . Clearly, the symmetries (1) contain determinism killing symmetries.

It is also worth noting that if the structure of spacetime becomes very minimal, no interesting laws of motion, deterministic or not, seem possible. For example, suppose that the time metric and the space metric are stripped from Leibnizian spacetime, leaving only the planes of absolute simultaneity. And suppose that the laws of physics specify that the world is filled with a plenum of constant mass dust particles and that the world lines of these particles are smooth curves that never cross. Then either every smooth, non-crossing motion of the dust is allowed by the laws of motion or none is, for any two such motions are connected by a symmetry of this minimal spacetime.

Two different strategies for saving determinism in the face of the above construction can be tried. They correspond to radically different attitudes towards

¹⁷A diffeomorphism d of the manifold \mathcal{M} is a one-one mapping of \mathcal{M} onto itself that preserves \mathcal{M} 's differentiable structure. For the sake of concreteness, assume that d is C^∞ .

¹⁸For on the assumption that the laws are (say) differential equations relating the O_i and P_j , they cannot be sensitive to the “bare identity” of the points of \mathcal{M} at which the O_i and P_j take some given values. This diffeomorphism invariance of the laws is one of the ingredients of what is called substantive general covariance (see section 6.2). One might contemplate breaking diffeomorphism invariance by introducing names for individual spacetime points; but the occurrence of such names would violate the “universal” character that laws are supposed to have.

¹⁹The details of various classical spacetime structures are to be found in [Earman, 1989].

the ontology of spacetime. The first strategy is to beef up the structure of the background spacetime. Adding a standard of rotation kills the time dependence in $R_\beta^\alpha(t)$, producing what is called *Maxwellian spacetime*. But since the $a^\alpha(t)$ are still arbitrary functions of t there remain determinism killing symmetries. Adding a standard of inertial or straight line motion linearizes the $a^\alpha(t)$ to $v^\alpha t + c^\alpha$, where the v^α and c^α are constants, producing *neo-Newtonian spacetime*²⁰ whose symmetries are given by the familiar Galilean transformations

$$(2) \quad \begin{aligned} x^\alpha &\rightarrow x'^\alpha = R_\beta^\alpha x^\beta + v^\alpha t + c^\alpha & \alpha, \beta = 1, 2, 3. \\ t &\rightarrow t' = t + \text{const} \end{aligned}$$

The mappings indicated by (2) do not contain determinism killing symmetries since if such a map is the identity map for a finite stretch of time, no matter how short, then it is the identity map period. Note that this way of saving determinism carries with it an allegiance to “absolute” quantities of motion: in neo-Newtonian spacetime it makes good sense to ask whether an isolated particle is accelerating or whether an isolated extended body is rotating. To be sure, this absolute acceleration and rotation can be called ‘relational’ quantities, but the second place in the relation is provided by the structure of the spacetime — in particular, by the inertial structure — and not by other material bodies, as is contemplated by those who champion relational accounts of motion.

The second strategy for saving determinism proceeds not by beefing up the structure of the background spacetime but by attacking a hidden assumption of the above construction — the “container view” of spacetime. Picturesquely, this assumption amounts to thinking of spacetime as a medium in which particles and fields reside. More precisely, in terms of the above apparatus, it amounts to the assumption that $\langle \mathcal{M}, O_1, O_2, \dots, O_M, P_1, P_2, \dots, P_N \rangle$ and $\langle \mathcal{M}, O_1, O_2, \dots, O_M, d^*P_1, d^*P_2, \dots, d^*P_N \rangle$, where d is any diffeomorphism of \mathcal{M} such that $d^*P_j \neq P_j$ for some j , describe different physical situations, even when d is a spacetime symmetry, i.e. $d^*O_i = O_i$ for all i . Rejecting the container view leads to (one form of) relationism about spacetime. A spacetime relationist will take the above construction to show that, on pain of abandoning the possibility of determinism, those who are relationists about motion should also be relationists about spacetime. Relationists about motion hold that talk of absolute motion is nonsensical and that all meaningful talk about motion must be construed as talk about the relative motions of material bodies. They are, thus, unable to avail themselves of the beef-up strategy for saving determinism; so, if they want determinism, they must grasp the lifeline of relationism about spacetime.

Relationism about motion is a venerable position, but historically it has been characterized more by promises than performances. Newton produced a stunningly successful theory of the motions of terrestrial and celestial bodies. Newton’s opponents promised that they could produce theories just as empirically adequate

²⁰Full Newtonian spacetime adds a distinguished inertial frame — ‘absolute space’ — thus killing the velocity term in (2).

and as explanatorily powerful as his without resorting to the absolute quantities of motion he postulated. But mainly what they produced was bluster rather than workable theories.²¹ Only in the twentieth century were such theories constructed (see [Barbour, 1974] and [Barbour and Bertotti, 1977]; and see [Barbour, 1999] for the historical antecedents of these theories), well after Einstein's GTR swept away the notion of a fixed background spacetime and radically altered the terms of the absolute vs. relational debate.

3.3 *Determinism and gauge symmetries*

When philosophers hear the word "gauge" they think of elementary particle physics, Yang-Mills theories, etc. This is a myopic view. Examples of non-trivial gauge freedom arise even in classical physics — in fact, we just encountered an example in the preceding subsection. The gauge notion arises for a theory where there is "surplus structure" (to use Michael Redhead's phrase) in the sense that the state descriptions provided by the theory correspond many-one to physical states. For such a theory a gauge transformation is, by definition, a transformation that connects those descriptions that correspond to the same physical state.

The history of physics shows that the primary reason for seeing gauge freedom at work is to maintain determinism. This thesis has solid support for the class of cases of most relevance to modern physics, viz. where the equations of motion/field equations are derivable from an action principle and, thus, the equations of motion are in the form of Euler-Lagrange equations.²² When the Lagrangian is non-singular, the appropriate initial data picks out a unique solution of the Euler-Lagrange equations and Laplacian determinism holds.²³ If, however, the action admits as variational symmetries a Lie group whose parameters are arbitrary functions of the independent variables, then we have a case of underdetermination because Noether's second theorem tells us that the Euler-Lagrange equations have to satisfy a set of mathematical identities.²⁴ When these independent variables include time, arbitrary functions of time will show up in solutions to the Euler-Lagrange equations, apparently wrecking determinism.

The point can be illustrated with the help of a humble example of particle mechanics constructed within the Maxwellian spacetime introduced in the preceding subsection. An appropriate Lagrangian invariant under the symmetries of this spacetime is given by

$$(3) \quad L = \sum \sum_{j < k} \frac{m_j m_k}{2M} (\dot{\mathbf{x}}_j - \dot{\mathbf{x}}_k)^2 - V(|\mathbf{x}_j - \mathbf{x}_k|), \quad M := \sum_i m_i.$$

²¹Newton's opponents were correct in one respect: Newton's postulation of absolute space, in the sense of a distinguished inertial frame was not needed to support his laws of motion.

²²See [Butterfield, this vol.] and [Belot, this vol.] for accounts of the Lagrangian and Hamiltonian formalisms.

²³At least if the continuity assumptions discussed in Section 3.5 below are imposed.

²⁴For an account of the Noether theorems, see [Brading and Brown, 2003] and [Brading and Castellani, this vol.].

This Lagrangian is singular in the sense that Hessian matrix $\partial^2 L / \partial \dot{\mathbf{x}}_i \partial \dot{\mathbf{x}}_j$ does not have an inverse. The Euler-Lagrange equations are

$$(4) \quad \frac{d}{dt} \left(m_i \dot{\mathbf{x}}_j - \frac{1}{M} \sum_k m_k \dot{\mathbf{x}}_k \right) = \frac{\partial V}{\partial \dot{\mathbf{x}}_i}.$$

These equations do not determine the evolution of the particle positions uniquely: if $\mathbf{x}_i(t)$ is a solution, so is $\mathbf{x}'_i(t) = \mathbf{x}_i(t) + \mathbf{f}(t)$, for arbitrary $\mathbf{f}(t)$, confirming the intuitive argument given above for the apparent breakdown of determinism. Determinism can be restored by taking the transformation $\mathbf{x}_i(t) \rightarrow \mathbf{x}_i(t) + \mathbf{f}(t)$ as a gauge transformation.

The systematic development of this approach to gauge was carried out by P. A. M. Dirac in the context of the Hamiltonian formalism.²⁵ A singular Lagrangian system corresponds to a constrained Hamiltonian system. The *primary constraints* appear as a result of the definition of the canonical momenta. (In the simple case of a first-order Lagrangian $L(q, \dot{q}, t)$, where q stands for the configuration variables and $\dot{q} := dq/dt$, the canonical momentum is $p := \partial L / \partial \dot{q}$.) The *secondary constraints* arise as a consequence of the demand that the primary constraints be preserved by the motion. The total set of constraints picks out the *constraint surface* $\mathcal{C}(q, p)$ of the Hamiltonian phase space $\Gamma(q, p)$. The *first class constraints* are those that commute on $\mathcal{C}(q, p)$ with all of the constraints. It is these first class constraints that are taken as the generators of the gauge transformations. The gauge invariant quantities (a.k.a. “observables”) are then the phase function $F : \Gamma(q, p) \rightarrow \mathbb{R}$ that are constant along the gauge orbits.

Applying the formalism to our toy case of particle mechanics in Maxwellian spacetime, the canonical momenta are:

$$(5) \quad \mathbf{p}_i := \frac{\partial L}{\partial \dot{\mathbf{x}}_i} = \frac{m_i}{M} \sum_k m_k (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_k) = m_i \dot{\mathbf{x}}_i - \frac{m_i}{M} \sum_k m_k \dot{\mathbf{x}}_k.$$

These momenta are not independent but must satisfy three primary constraints, which require the vanishing of the x , y , and z -components of the total momentum:

$$(6) \quad \phi_\alpha = \sum_i p_i^\alpha = 0, \quad \alpha = 1, 2, 3.$$

These primary constraints are the only constraints — there are no secondary constraints — and they are all first class. These constraints generate in each configuration variable \mathbf{x}_i the same gauge freedom; namely, a Euclidean shift given by the same arbitrary function of time. The gauge invariant variables, such relative particle positions and relative particle momenta, do evolve deterministically.

The technical elaboration of the constraint formalism is complicated, but one should not lose sight of the fact that the desire to save determinism is the motivation driving the enterprise. Here is a relevant passage from [Henneaux and

²⁵The standard reference on these matters is [Henneaux and Teitelboim, 1992]. For a user friendly treatment of this formalism, see [Earman, 2003].

Teitelboim, 1992], one of the standard references on constrained Hamiltonian systems:

The presence of arbitrary functions ... in the total Hamiltonian tells us that not all the q 's and p 's [the configuration variables and their canonical momenta] are observable [i.e. genuine physical magnitudes]. In other words, although the physical state is uniquely defined once a set of q 's and p 's is given, the converse is not true — i.e., there is more than one set of values of the canonical variables representing a given physical state. To see how this conclusion comes about, we note that if we are given an initial set of canonical variables at the time t_1 and thereby completely define the physical state at that time, we expect the equations of motion to *fully determine the physical state at other times*. Thus, by definition, any ambiguity in the value of the canonical variables at $t_2 \neq t_1$ should be a physically irrelevant ambiguity. [pp. 16–17]

As suggested by the quotation, the standard reaction to the apparent failure of determinism is to blame the appearance on the redundancy of the descriptive apparatus: the correspondence between the state descriptions in terms of the original variables — the q 's and p 's — and the physical state is many-to-one; when this descriptive redundancy is removed, the physical state is seen to evolve deterministically. There may be technical difficulties in carrying through this reaction. For example, attempting to produce a reduced phase space — whose state descriptions corresponding one-one to physical states — by quotienting out the gauge orbits can result in singularities. But when such technical obstructions are not met, normal (i.e. unconstrained) Hamiltonian dynamics applies to the reduced phase space, and the reduced phase space variables evolve deterministically.

In addition to this standard reaction to the apparent failure of determinism in the above examples, two others are possible. The first heterodoxy takes the apparent violation of determinism to be genuine. This amounts to (a) treating what the constraint formalism counts as gauge dependent quantities as genuine physical magnitudes, and (b) denying that these magnitudes are governed by laws which, when conjoined with the laws already in play, restore determinism. The second heterodoxy accepts the orthodox conclusion that the apparent failure of determinism is merely apparent; but it departs from orthodoxy by accepting (a), and it departs from the first heterodoxy by denying (b) and, accordingly, postulates the existence of additional laws that restore determinism. Instances that superficially conform to part (a) of the two heterodoxies are easy to construct from examples found in physics texts where the initial value problem is solved by supplementing the equations of motion, stated in terms of gauge-dependent variables, with a gauge condition that fixes a unique solution. For instance, Maxwell's equations written in terms of electromagnetic potentials do not determine a unique solution corresponding to the initial values of the potentials and their time derivatives. Imposing the Lorentz gauge condition converts Maxwell's equations to second or-

der hyperbolic partial differential equations (pdes) that do admit an initial value formulation (see Section 4.2).²⁶ Similar examples can be concocted in general relativity theory where orthodoxy treats the metric potentials as gauge variables (see Section 6.2). In these examples orthodoxy is aiming to get at the values of the gauge independent variables via a choice of gauge. If this aim is not kept clearly in mind, the procedure creates the illusion that gauge-dependent variables have physical significance. It is exactly this illusion that the two heterodoxies take as real. The second heterodoxy amounts to taking the gauge conditions not as matters of calculational convenience but as additional physical laws. I know of no historical examples where this heterodoxy has led to fruitful developments in physics.

Since there is no *a priori* guarantee that determinism is true, the fact that the orthodox reading of the constraint formalism guarantees that the equations of motion admit an initial value formulation must mean that substantive assumptions that favor determinism are built into the formalism. That is indeed the case, for the Lagrangian/Hamiltonian formalism imposes a structure on the space of solutions: in the geometric language explained in Chapters 1 and 2 of this volume, the space of solutions has a symplectic or pre-symplectic structure. This formalism certainly is not guaranteed to be applicable to all of the equations of motion the Creator might have chosen as laws of motion; indeed, it is not even guaranteed to be applicable to all Newtonian type second order odes. In the 1880s Helmholtz found a set of necessary conditions for equations of this type to be derivable from an action principle; these conditions were later proved to be (locally) sufficient as well as necessary. After more than a century, the problem of finding necessary and sufficient conditions for more general types of equations of motion, whether in the form of odes or pdes, to be derivable from an action principle is still an active research topic.²⁷

3.4 *Determinism for fields and fluids in Newtonian physics*

Newtonian gravitational theory can be construed as a field theory. The gravitational force is given by $\mathbf{F}_{grav} = -\nabla\varphi$, where the gravitational potential φ satisfies the Poisson equation

$$(7) \quad \nabla^2\varphi = \rho$$

with ρ being the mass density. If φ is a solution to Poisson's equation, then so is $\varphi' = \varphi + g(\mathbf{x})f(t)$ where $g(\mathbf{x})$ is a linear function of the spatial variables and $f(t)$

²⁶Where \mathbf{A} is the vector potential and Φ is the scalar potential, the Lorentz gauge requires that

$$\nabla \cdot \mathbf{A} + \frac{\partial \Phi}{\partial t} = 0$$

(with the velocity of light set to unity).

²⁷Mathematicians discuss this issue under the heading of the "inverse problem." For precise formulations of the problem and surveys of results, see [Anderson and Thompson, 1992] and [Prince, 2000].

is an arbitrary function of t . Choose f so that $f(t) = 0$ for $t \leq 0$ but $f(t) > 0$ for $t > 0$. The extra gravitational force, proportional to $f(t)$, that a test particle experiences in the primed solution after $t = 0$ is undetermined by anything in the past.

The determinism wrecking solutions to (7) can be ruled out by demanding that gravitational forces be tied to sources. But to dismiss homogeneous solutions to the Poisson equation is to move in the direction of treating the Newtonian gravitational field as a mere mathematical device that is useful in describing gravitational interactions which, at base, are really direct particle interactions.²⁸ In this way determinism helps to settle the ontology of Newtonian physics: the insistence on determinism in Newtonian physics demotes fields to second-class status. In relativistic physics fields come into their own, and one of the reasons is that the relativistic spacetime structure supports field equations that guarantee deterministic evolution of the fields (see Section 4.2).

In the Newtonian setting the field equations that naturally arise are elliptic (e.g. the Poisson equation) or parabolic, and neither type supports determinism-without-crutches. An example of the latter type of equation is the classical heat equation

$$(8) \quad \nabla^2 \Phi = \kappa \frac{\partial \Phi}{\partial t}$$

where Φ is the temperature variable and κ is the coefficient of heat conductivity.²⁹ Solutions to (8) can cease to exist after a finite time because the temperature “blows up.” Uniqueness also fails since, using the fact that the heat equation propagates heat arbitrarily fast, it is possible to construct surprise solutions Φ_s with the properties that (i) Φ_s is infinitely differentiable, and (ii) $\Phi_s(\mathbf{x}, t) = 0$ for all $t \leq 0$ but $\Phi_s(\mathbf{x}, t) \neq 0$ for $t > 0$ (see [John, 1982, Sec. 7.1]). Because (8) is linear, if Φ is a solution then so is $\Phi' = \Phi + \Phi_s$. And since Φ and Φ' agree for all $t \leq 0$ but differ for $t > 0$, the existence of the surprise solutions completely wrecks determinism.

Uniqueness of solution to (8) can be restored by adding the requirement that $\Phi \geq 0$, as befits its intended interpretation of Φ as temperature; for Widder [1975, 157] has shown that if a solution of $\Phi(\mathbf{x}, t)$ of (8) vanishes at $t = 0$ and is non-negative for all \mathbf{x} and all $t \geq 0$, then it must be identically zero. But one could have wished that, rather than having to use a stipulation of non-negativity to shore up determinism, determinism could be established and then used to show that if the temperature distribution at $t = 0$ is non-negative for all \mathbf{x} , then the uniquely determined evolution keeps the temperature non-negative. Alternatively, both uniqueness and existence of solutions of (8) can be obtained by limiting the

²⁸This demotion of the status of the Newtonian gravitational field can also be supported by the fact that, unlike the fields that will be encountered in relativistic theories, it carries no energy or momentum.

²⁹The fact that this equation is not Galilean invariant need cause no concern since Φ implicitly refers to the temperature of a medium whose rest frame is the preferred frame for describing heat diffusion.

growth of $|\Phi(\mathbf{x}, t)|$ as $|\mathbf{x}| \rightarrow \infty$. But again one could have wished that such limits on growth could be derived as a consequence of the deterministic evolution rather than having to be stipulated as conditions that enable determinism.

Appearances of begging the question in favor of determinism could be avoided by providing at the outset a clear distinction between kinematics and dynamics, the former being a specification of the space of possible states. For example, a limit on the growth of quantum mechanical wave functions does not beg the question of determinism provided by the Schrödinger equation since the limit follows from the condition that the wave function is an element of a Hilbert space, which is part of the kinematical prescription of QM (see Section 5). Since this prescription is concocted to underwrite the probability interpretation of the wave function, we get the ironic result that the introduction of probabilities, which seems to doom determinism, also serves to support it. The example immediately above, as well as the examples of the preceding subsection and the one at the beginning of this subsection, indicate that in classical physics the kinematical/dynamical distinction can sometimes be relatively fluid and that considerations of determinism are used in deciding where to draw the line. The following example will reinforce this moral.³⁰

The Navier-Stokes equations for an incompressible fluid moving in \mathbb{R}^N , $N = 2, 3$, read

$$\frac{D\mathbf{u}}{dt} = -\nabla p + v\Delta\mathbf{u} \quad (9a)$$

$$\operatorname{div}(\mathbf{u}) = 0 \quad (9b)$$

where $\mathbf{u}(\mathbf{x}, t) = (u^1, u^2, \dots, u^N)$ is the velocity of the fluid, $p(\mathbf{x}, t)$ is the pressure, $v = \text{const.} \geq 0$ is the coefficient of viscosity, and $D/dt := \partial/\partial t + \sum_{j=1}^N u^j \partial/\partial x^j$ is

the convective derivative (see Foias et al. 2001 for a comprehensive survey). If the fluid is subject to an external force, an extra term has to be added to the right hand side of (9a). The Euler equations are the special case where $v = 0$. The initial value problem for (9a-b) is posed by giving the initial data

$$(9) \quad \mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x})$$

³⁰Another reaction to the problems of determinism posed by (8) is to postpone them on the grounds that (8) is merely a phenomenological equation; heat is molecular motion and, thus, the fate of determinism ultimately rests with the character of laws of particle motion. It will be seen below, however, that in order to guarantee determinism for particle motion the helping hand of the stipulation of boundary conditions at infinity is sometimes needed. In any case, the postponement strategy taken to its logical conclusion would mean that no judgment about determinism would be forthcoming until we are in possession of the final theory-of-everything. It seems a better strategy to do today the philosophy of today's physics while recognizing, of course, that today's best theory may be superseded by a better future theory that delivers a different message about determinism.

where $\mathbf{u}_0(\mathbf{x})$ is a smooth (C^∞) divergence-free vector field, and is solved by smooth functions $\mathbf{u}, p \in C^\infty(\mathbb{R}^N \times [0, \infty))$ satisfying (9)-(10). For physically reasonable solutions it is required both that $\mathbf{u}_0(\mathbf{x})$ should not grow too large as $|\mathbf{x}| \rightarrow \infty$ and that the energy of the fluid is bounded for all time:

$$(10) \quad \int_{\mathbb{R}^N} |\mathbf{u}(\mathbf{x}, t)|^2 dx < \infty \text{ for all } t > 0.$$

When $\nu = 0$ the energy is conserved, whereas for $\nu > 0$ it dissipates.

For $N = 2$ it is known that a physically reasonable smooth solution exists for any given $\mathbf{u}_0(\mathbf{x})$. For $N = 3$ the problem is open. However, for this case it is known that the problem has a positive solution if the time interval $[0, \infty)$ for which the solution is required to exist is replaced by $[0, T)$ where T is a possibly finite number that depends on $\mathbf{u}_0(\mathbf{x})$. When T is finite it is known as the “blowup time” since $|\mathbf{u}(\mathbf{x}, t)|$ must become unbounded as t approaches T . For the Euler equations a finite blowup time implies that the vorticity (i.e. the *curl* of $\mathbf{u}(\mathbf{x}, t)$) becomes unbounded as t approaches T .

Smooth solutions to the Navier-Stokes equations, when they exist, are known to be unique. This claim would seem to be belied by the symmetries of the Navier-Stokes equations since if $\mathbf{u}(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t)$, $p(\mathbf{x}, t) = g(\mathbf{x}, t)$ is a solution then so is the transformed $\tilde{\mathbf{u}}(\mathbf{x}, t) = \mathbf{f}(\mathbf{x} - \varepsilon\alpha(t), t) + \varepsilon\alpha_t$, $\tilde{p}(\mathbf{x}, t) = g(\mathbf{x} - \varepsilon\alpha(t), t) - \varepsilon\mathbf{x} \cdot \alpha_t + \frac{1}{2}\varepsilon^2\alpha_{tt}$, where $\alpha(t)$ is an arbitrary smooth function of t alone (see Olver 1993, pp. 130 and 177 (Exer. 2.15)). Choosing $\alpha(t)$ such that $\alpha(0) = \alpha_t(0) = \alpha_{tt}(0) = 0$ but $\alpha(t) \neq 0$ for $t > 0$ results in different solutions for the same initial data unless $\mathbf{f}(\mathbf{x} - \varepsilon\alpha(t), t) + \varepsilon\alpha_t = \mathbf{f}(\mathbf{x}, t)$. However, the transformed solution violates the finiteness of energy condition (11).

The situation on the existence of solutions can be improved as follows. Multiplying (9a-b) by a smooth test function and integrating by parts over \mathbf{x} and t produces integral equations that are well-defined for any $\mathbf{u}(\mathbf{x}, t)$ and $p(\mathbf{x}, t)$ that are respectively L^2 (square integrable) and L^1 (integrable). Such a pair is called a *weak solution* if it satisfies the integral equations for all test functions. Moving from smooth to weak solutions permits the proof of the existence of a solution for all time. But the move reopens the issue of uniqueness, for the uniqueness of weak solutions for the Navier-Stokes equations is not settled. A striking non-uniqueness result for weak solutions of the Euler equations comes from the construction by Scheffer [1994] and Shnirelman [1997] of self-exciting/self-destroying weak solutions: $\mathbf{u}(\mathbf{x}, t) \equiv 0$ for $t < -1$ and $t > 1$, but is non-zero between these times in a compact region of \mathbb{R}^3 .

It is remarkable that basic questions about determinism for classical equations of motion remain unsettled and that these questions turn on issues that mathematicians regard as worthy of attention. Settling the existence question for smooth solutions for the Navier-Stokes equations in the case of $N = 3$ brings a \$1 million award from the Clay Mathematics Institute (see [Fefferman, 2000]).

3.5 Continuity issues

Consider a single particle of mass m moving on the real line \mathbb{R} in a potential $V(x)$, $x \in \mathbb{R}$. The standard existence and uniqueness theorems for the initial value problem of odes can be used to show that the Newtonian equation of motion

$$(11) \quad m\ddot{x} = F(x) := -\frac{dV}{dx}$$

has a locally (in time) unique solution if the force function $F(x)$ satisfies a Lipschitz condition.³¹ An example of a potential that violates the Lipschitz condition at the origin is $-\frac{9}{2}|x|^{4/3}$. For the initial data $x(0) = 0 = \dot{x}(0)$ there are multiple solutions of (12): $x(t) \equiv 0$, $x(t) = t^3$, and $x(t) = -t^3$, where m has been set to unity for convenience. In addition, there are also solutions $x(t)$ where $x(t) = 0$ for $t < k$ and $\pm(t-k)^3$ for $t \geq k$, where k is any positive constant. That such force functions do not turn up in realistic physical situations is an indication that Nature has some respect for determinism. In QM it turns out that Nature can respect determinism while accommodating some of the non-Lipschitz potentials that would wreck Newtonian determinism (see Section 5.2).

3.6 The breakdown of classical solutions

Consider again the case of a single particle of mass m moving on the real line \mathbb{R} in a potential $V(x)$, and suppose that $V(x)$ satisfies the Lipschitz condition, guaranteeing a temporally local unique solution for the initial value problem for the Newtonian equations of motion. However, determinism can fail if the potential is such that the particle is accelerated off to $-\infty$ or $+\infty$ in a finite time.³² Past determinism is violated because two such solutions can agree for all future times $t \geq t^*$ (say) — no particle is present at these times anywhere in space — but disagree at past times $t < t^*$ on the position and/or velocity of the particle when it is present in space. Since the potential is assumed to be time independent, the equations of motion are time reversal invariant, so taking the time reverses of these escape solutions produces solutions in which hitherto empty space is invaded by particles appearing from spatial infinity. These invader solutions provide violations of future determinism. Piecing together escape and invader solutions produces further insults to determinism.

In the 1890's Paul Painlevé conjectured that for $N > 3$ point mass particles moving in \mathbb{R}^3 under their mutually attractive Newtonian gravitational forces, there exist solutions to the Newtonian equations of motion exhibiting non-collision singularities, i.e. although the particles do not collide, the solution ceases to exist

³¹ $F(x)$ satisfies the Lipschitz condition in an interval $(a, b) \subset \mathbb{R}$ if there is a constant $K > 0$ such that $|F(x_1) - F(x_2)| \leq K|x_1 - x_2|$ for all $x_1, x_2 \in (a, b)$. A sufficient condition for this is that dF/dx exists, is continuous, and $|dF/dx| \leq K$ on (a, b) for some $K > 0$.

³²See [Reed and Simon, 1975, Theorem X.5] for necessary and sufficient conditions for this to happen.

after a finite time. Hugo von Zeipel [1908] showed that in such a solution the particle positions must become unbounded in a finite time. Finally, near the close of the 20th century Xia [1992] proved Painlevé conjecture by showing that for $N = 5$ point mass particles, the Newtonian equations of motion admit solutions in which the particles do not collide but nevertheless manage to accelerate themselves off to spatial infinity in a finite time (see [Saari and Xia, 1995] for an accessible survey).

Determinism can recoup its fortunes by means of the device, already mentioned above, of supplementing the usual initial conditions with boundary conditions at infinity. Or consolation can be taken from two remarks. The first remark is that in the natural phase space measure, the set of initial conditions that lead to Xia type escape solutions has measure zero. But it is unknown whether the same is true of all non-collision singularities. The second remark is that the non-collision singularities result from the unrealistic idealization of point mass particles that can achieve unbounded velocities in a finite time by drawing on an infinitely deep potential well. This remark does not suffice to save determinism when an infinity of finite sized particles are considered, as we will see in the next subsection.

It is interesting to note that for point particles moving under mutually attractive Newtonian gravitational forces, QM cures both the collision³³ and non-collision singularities that can spell the breakdown of classical solutions (see Section 5.2). This is more than a mere mathematical curiosity since it is an important ingredient in the explanation of the existence and stability of the hydrogen atom.

3.7 *Infinite collections*

Consider a collection of billiard balls confined to move along a straight line in Euclidean space. Suppose that the balls act only by contact, that only binary collisions occur, and that each such collision obeys the classical laws of elastic impact. Surely, the reader will say, such a system is as deterministic as it gets. This is so, *if* the collection is finite. But if the collection is infinite and unbounded velocities are permitted, then determinism fails because even with all of the announced restrictions in place the system can seemingly self-excite itself (see [Lanford, 1974]). Pérez Laraudogoitia [2001] shows how to use such infinite collections to create an analogue of the escape solution of the preceding subsection where all of the particles disappear in a finite amount of time. The time reverse of this scenario is one in which space is initially empty, and then without warning an infinite stream of billiard balls pour in from spatial infinity.

Legislating against unbounded velocities or imposing boundary conditions at infinity does not suffice to restore determinism if the billiard balls can be made arbitrarily small [Pérez Laraudogoitia, 2001]. For then a countably infinite collection of them can be Zeno packed into a finite spatial interval, say $(0, 1]$, by placing the center of the first ball at 1, the second at $1/2$, the third at $1/4$, etc. Assume for ease of illustration that all the balls have equal mass ($\equiv 1$). A unit mass cue

³³A collision singularity occurs when two or more of the point particles collide and the solution cannot be continued through the collision time.

ball moving with unit speed from right to left collides with the first ball and sends a ripple through the Zeno string that lasts for unit time, at the end of which all of the balls are at rest. The boring history in which all the balls are at rest for all time is, of course, also a solution of the laws of impact. Comparing this boring history with the previous one shows that past Laplacian determinism is violated.³⁴

This failure of determinism carries with it a violation of the conservation and energy momentum, albeit in a weak sense; namely, in the inertial frame in which the object balls are initially at rest, the total energy and the total momentum each have different values before and after the collisions start, but in every other inertial frame there is no violation simply because the values are infinite both before and after the collisions.³⁵ Pérez Laraudogoitia [2005] has shown how to construct scenarios in which there is a strong violation of conservation of energy and momentum in that the violation occurs in every inertial frame.

3.8 Domains of dependence

With some artificiality one of the threats to classical determinism discussed above can be summarized using a concept that will also prove very helpful in comparing the fortunes of determinism in classical physics and in relativistic physics. By a *causal curve* let us understand a (piecewise) smooth curve in spacetime that represents the spacetime trajectory for a physically possible transfer of energy/momentum. Define the *future domain of dependence*, $D^+(S)$, of a spacetime region S as the set of all spacetime points p such that any past directed causal curve with future endpoint at p and no past endpoint intersects S . The *past domain of dependence* $D^-(S)$ of S is defined analogously. And the *total domain of dependence* $D(S)$ is the union $D^+(S) \cup D^-(S)$. If $p \notin D(S)$ then it would seem that the state in region S does not suffice to determine the state at p since there is a possible causal process that passes through p but never registers on S .

Since neither the kinematics nor the dynamics of classical physics place an upper bound on the velocity at which energy/momentum can be transferred, it would seem that in principle any timelike curve — i.e. any (piecewise) smooth curve oblique to the planes of absolute simultaneity — can count as a causal curve, and as a consequence $D(S) = \emptyset$ even when S is taken to be an entire plane of absolute simultaneity. The examples from Sections 3.4, 3.6, and 3.7 show how the “in principle” can be realized by some systems satisfying Newtonian laws of motion.

We have seen that some threats to classical determinism can be met by beefing up the structure of classical spacetime. And so it is with the threat currently under consideration. *Full Newtonian spacetime* is what results from neo-Newtonian

³⁴The time reverse of the interesting history starts with all the balls initially at rest, and then subsequently the collection self-excites, sending a ripple of collisions from left to right and ejecting the cue ball. If this self-exciting history is physically possible, then future laplacian determinism is violated. However, it might be rejected on the grounds that it violated Newton’s first law of motion.

³⁵For a comment on how the availability of an infinite amount of momentum/energy renders the indeterminism unsurprising, see [Norton, 1999, 1268].

spacetime by adding absolute space in the form of a distinguished inertial frame ('absolute space'). In this setting the spacetime symmetries are small enough that there are now finite invariant velocities (intuitively, velocities as measured relative to absolute space), and thus laws can be formulated that set a finite upper bound on the absolute velocity of causal propagation. Nor is this move necessarily *ad hoc* as shown, for example, by the fact that the formulation of Maxwell's laws of electromagnetism in a classical spacetime setting evidently requires the services of a distinguished inertial frame, the velocity of light c being the velocity as measured in this frame.

But, as is well known, such a formulation is embarrassed by the undetectability of motion with respect to absolute space. This embarrassment provides a direct (albeit anachronistic) route from classical to relativistic spacetime. Adopting for classical spacetimes the same geometric language used in the special and general theories of relativity (see [Earman, 1989, Ch. 2]), absolute space is represented by a covariantly constant timelike vector field A^a , the integral curves of which are the world lines of the points of absolute space. The space metric is represented by a degenerate second rank contravariant tensor h^{ab} , which together with A^a defines a tensor that is formally a Minkowski metric: $\eta^{ab} := h^{ab} - A^a A^b$. The unobservability of absolute motion means that there is no preferred way to split η^{ab} into an h^{ab} part and a $A^a A^b$ part, suggesting that η^{ab} is physically as well as formally a Lorentz metric. As we will see in Section 4.1, this puts determinism on much firmer ground in that domains of dependence of local or global time slices are non-empty in the spacetime setting of STR.

3.9 Determinism, predictability, and chaos

Laplace's vision of a deterministic universe makes reference to an "intelligence" (which commentators have dubbed 'Laplace's Demon'):

We ought to regard the present state of the universe as the effect of its antecedent state and as the cause of the state that is to follow. An intelligence knowing all of the forces acting in nature at a given instant, as well as the momentary positions of all things in the universe, would be able to comprehend in one single formula the motions of the largest bodies as well as the lightest atoms in the world, provided that its intellect were sufficiently powerful to subject all data to analysis; to it nothing would be uncertain, the future as well as the past would be present to its eyes.³⁶

³⁶[Laplace, 1820]. English translation from [Nagel, 1961, 281-282]. More than a century earlier Leibniz espoused a similar view: "[O]ne sees then that everything proceeds mathematically — that is, infallibly — in the whole wide world, so that if someone could have sufficient insight into the inner parts of things, and in addition has remembrance and intelligence enough to consider all the circumstances and to take them into account, he would be a prophet and would see the future in the present as in a mirror." Quoted from [Cassirer, 1956, 12].

Perhaps by taking Laplace's vision too literally, philosophers and physicists alike conflate determinism and predictability. The conflation leads them to reason as follows: here is a case where predictability fails; thus, here is a case where determinism fails.³⁷ This is a mistake that derives from a failure to distinguish determinism — an ontological doctrine about how the world evolves — from predictability — an epistemic doctrine about what can be inferred, by various restricted means, about the future (or past) state of the world from a knowledge of its present state.

There is, however, an interesting connection between determinism and practical predictability for laws of motion that admit an initial value problem that is *well-posed* in the sense that, in some appropriate topology, the solutions depend continuously on the initial data.³⁸ The standard existence and uniqueness proofs for the initial value problem for the odes used in particle mechanics also furnish a proof of well-posedness, which can be traced to the fact that the existence proof is constructive in that it gives a procedure for constructing a series of approximations that converge to the solution determined by the initial data.

To illustrate the implications of well-posedness for predictability, consider the toy case of a system consisting of a single massive particle obeying Newtonian equations of motion. If a suitable Lipschitz condition is satisfied, then for any given values of the position $q(0)$ and velocity $\dot{q}(0)$ of the particle at $t = 0$ there exists (for some finite time interval surrounding $t = 0$) a unique solution: symbolically $q(t) = F(q(0), \dot{q}(0), t)$. And further, since this initial value problem is well-posed, for any fixed $t > 0$ (within the interval for which the solution is guaranteed to exist), F is a continuous function of $q(0)$ and $\dot{q}(0)$. Suppose then that the practical prediction task is to forecast the actual position $\bar{q}(t^*)$ of the particle at some given $t^* > 0$ with an accuracy of $\epsilon > 0$, and suppose that although measurements of position or velocity are not error free, the errors can be made arbitrarily small. By the continuity of F , there exist $\delta_1 > 0$ and $\delta_2 > 0$ such that if $|q(0) - \bar{q}(0)| < \delta_1$ and $|\dot{q}(0) - \bar{\dot{q}}(0)| < \delta_2$, then $|q(t^*) - \bar{q}(t^*)| < \epsilon$. Thus, measuring at $t = 0$ the actual particle position and velocity with accuracies $\pm\delta_1/2$ and $\pm\delta_2/2$ respectively ensures that when the measured values are plugged into F , the value of the function for $t = t^*$ answers to the assigned prediction task. (Note, however, that since the actual initial state is unknown, so are the required accuracies $\pm\delta_1/2$ and $\pm\delta_2/2$, which may depend on the unknown state as well as on ϵ and t^* . This hitch could be overcome if there were minimum but non-zero values of δ_1 and δ_2 that

³⁷On the philosophical side, Karl Popper is the prime example. Popper [1982] goes so far as to formulate the doctrine of "scientific determinism" in terms of prediction tasks. An example on the physics side is Reichl [1992]: "[W]e now know that the assumption that Newton's equations are deterministic is a fallacy! Newton's equations are, of course, the correct starting point of mechanics, but in general they only allow us to determine [read: predict] the long time behavior of *integrable* mechanical systems, few of which can be found in nature" (pp. 2–3). I am happy to say that in the second edition of Reichl's book this passage is changed to "[W]e now know that the assumption that Newton's equations can *predict* the future is a fallacy!" [Reichl 2004, 3; italics added].

³⁸When the topology is that induced by a norm $\|\cdot\|$ on the instantaneous states represented by a function $s(t)$ of time, well-posedness requires that there is a non-decreasing, nonnegative function $C(t)$ such that $\|s(t)\| \leq C(t)\|s(0)\|$, $t > 0$, for any solution $s(t)$.

answered to the given prediction task whatever the initial state; but there is no *a priori* guarantee that such minimum values exist. A prior measurement with known accuracy of the position and velocity at some $t^{**} < 0$ will put bounds, which can be calculated from F , on the position and velocity at $t = 0$. And then the minimum values can be calculated for accuracies δ_1 and δ_2 of measurements at $t = 0$ that suffice for the required prediction task for any values of the position and velocity within the calculated bounds.)

Jacques Hadamard, who made seminal contributions to the Cauchy or initial value problem for pdes, took the terminology of “well-posed” (a.k.a. “properly posed”) quite literally. For he took it as a criterion for the proper mathematical description of a physical system that the equations of motion admit an initial value formulation in which the solution depends continuously on the initial data (see [Hadamard, 1923, 32]). However, the standard Courant-Hilbert reference work, *Methods of Mathematical Physics*, opines that

“properly posed” problems are by far not the only ones which appropriately reflect real phenomena. So far, unfortunately, little mathematical progress has been made in the important task of solving or even identifying such problems that are not “properly posed” but still are important and motivated by realistic situations. [1962, Vol. 2, 230].

Some progress can be found in [Payne, 1975] and the references cited therein.

Hadamard was of the opinion that if the time development of a system failed to depend continuously on the initial conditions, then “it would appear to us as being governed by pure chance (which, since Poincaré,³⁹ has been known to consist precisely in such a discontinuity in determinism) and not obeying any law whatever” [1923, 38]. Currently the opinion is that the appearance of chance in classical systems is due not to the failure of well-posedness but to the presence of chaos.

The introduction of *deterministic chaos* does not change any of the above conclusions about determinism and predictability. There is no generally agreed upon definition of chaos, but the target class of cases can be picked out either in terms of cause or effects. The cause is sensitive dependence of solutions on initial conditions, as indicated, for example, by positive Lyapunov exponents. The effects are various higher order ergodic properties, such as being a mixing system, being a K-system, being a Bernoulli system, etc.⁴⁰ Generally a sensitive dependence on initial conditions *plus* compactness of the state space is sufficient to secure such properties. The sensitive dependence of initial condition that is the root cause of chaotic behavior does not contradict the continuous dependence of solutions on initial data, and, therefore, does not undermine the task of predicting with any desired finite accuracy the state at a *fixed* future time, assuming that error in measuring the initial conditions can be made arbitrarily small. If, however, there

³⁹See Poincaré’s essay “Chance” in *Science and Method* [1952].

⁴⁰See Uffink, this volume, section 6.2, or [Lichtenberg and Leiberman, 1991] for definitions of these concepts.

is a fixed lower bound on the accuracy of measurements — say, because the measuring instruments are macroscopic and cannot make discriminations below some natural macroscopic scale — then the presence of deterministic chaos can make some prediction tasks impossible. In addition, the presence of chaos means that no matter how small the error (if non zero) in ascertaining the initial conditions, the accuracy with which the future state can be forecast degrades rapidly with time. To ensure the ability to predict with some given accuracy $\epsilon > 0$ for all $t > 0$ by ascertaining the initial conditions at $t = 0$ with sufficiently small error $\delta > 0$, it would be necessary to require not only well-posedness but *stability*, which is incompatible with chaos.⁴¹

Cases of classical chaos also show that determinism on the microlevel is not only compatible with stochastic behavior at the macro level but also that the deterministic microdynamics can ground the macro-stochasticity. For instance, the lowest order ergodic property — ergodicity — arguably justifies the use of the microcanonical probability distribution and provides for a relative frequency interpretation; for it implies that the microcanonical distribution is the only stationary distribution absolutely continuous with respect to Lebesgue measure and that the measure of a phase volume is equal to the limiting relative frequency of the time the phase point spends in the volume. In these cases there does not seem to be a valid contrast between “objective” and “epistemic” probabilities. The probabilities are epistemic in the sense that conditionalizing on a mathematically precise knowledge of the initial state reduces the outcome probability to 0 or 1. But the probabilities are not merely epistemic in the sense of merely expressing our ignorance, for they are supervenient on the underlying microdynamics.

Patrick Suppes [1991; 1993] has used such cases to argue that, because we are confined to the macrolevel, determinism becomes for us a “transcendental” issue since we cannot tell whether we are dealing with a case of irreducible stochasticity or a case of deterministic chaos. Although I feel some force to the argument, I am not entirely persuaded. There are two competing hypotheses to explain observed macro-stochasticity: it is due to micro-determinism plus sensitive dependence on initial conditions vs. it is due to irreducible micro-stochasticity. The work in recent decades on deterministic chaos supplies the details on how the first hypothesis can be implemented. The details of the second hypothesis need to be filled in; particular, it has to be explained how the observed macro-stochasticity supervenes on the postulated micro-stochasticity.⁴² And then it has to be demonstrated that the two hypotheses are underdetermined by all possible observations on the macrolevel. If both of these demands were met, we would be faced with a particular instance of the general challenge to scientific realism posed by underdetermination of theory by observational evidence, and all of the well-rehearsed moves and countermoves in the realism debate would come into play. But it is futile to fight these battles until some concrete version of the second hypothesis is presented.

⁴¹Stability with respect to a norm on states $s(t)$ requires that there is a constant C such that $\|s(t)\| \leq C\|s(0)\|$, $t > 0$, for any solution $s(t)$. Compare to footnote 38.

⁴²It is not obvious that micro-stochasticity always percolates up to the macro-level.

3.10 Laplacian demons, prediction, and computability

Since we are free to imagine demons with whatever powers we like, let us suppose that Laplace's Demon can ascertain the initial conditions of the system of interest with absolute mathematical precision. As for computational ability, let us suppose that the Demon has at its disposal a universal Turing machine. As impressive as these abilities are, they may not enable the Demon to predict the future state of the system even if it is deterministic. Returning to the example of the Newtonian particle from the preceding subsection, if the values of the position and velocity of the particle at time $t = 0$ are plugged into the function $F(q(0), \dot{q}(0), t)$ that specifies the solution $q(t)$, the result is a function $\mathcal{F}(t)$ of t ; and plugging different values of the initial conditions results in different $\mathcal{F}(t)$ — indeed, by the assumption of determinism, the $\mathcal{F}(t)$'s corresponding to different initial conditions must differ on any finite interval of time no matter how small. Since there is a continuum of distinct initial conditions, there is thus a continuum of distinct $\mathcal{F}(t)$'s. But only a countable number of these $\mathcal{F}(t)$'s will be Turing computable functions.⁴³ Thus, for most of the initial conditions the Demon encounters, it is unable to predict the corresponding particle position $q(t)$ at $t > 0$ by using its universal Turing machine to compute the value of $\mathcal{F}(t)$ at the relevant value of t — in Pitowsky's [1996] happy turn of phrase, the Demon must consult an Oracle in order to make a sure fire prediction.

However, if $q(0)$ and $\dot{q}(0)$ are both Turing computable real numbers, then an Oracle need not be consulted since the corresponding $\mathcal{F}(t)$ is a Turing computable function; and if t is a Turing computable real number, then so is $\mathcal{F}(t)$. This follows from the fact that the existence and uniqueness proofs for odes gives an effective procedure for generating a series of approximations that converges effectively to the solution; hence, if computable initial data are fed into the procedure, the result is an effectively computable solution function. Analogous results need not hold when the equations of motion are pdes. Jumping ahead to the relativistic context, the wave equation for a scalar field provides an example where Turing computability of initial conditions is not preserved by deterministic evolution (see Section 4.4).

A more interesting example where our version of Laplace's Demon must consult an Oracle has been discussed by Moore [1990; 1991] and Pitowsky [1996]. Moore constructed an embedding of an abstract universal Turing machine into a concrete classical mechanical system consisting of a particle bouncing between parabolic and flat mirrors arranged so that the motion of the particle is confined to a unit

⁴³The familiar notion of a Turing computable or recursive function is formulated for functions of the natural numbers, but it can be generalized so as to apply to functions of the real numbers. First, a computable real number x is defined as a limit of a computable sequence $\{r_n\}$ of rationals that converges effectively, i.e. there is a recursive function $f(n)$ such that $k \geq f(n)$ entails $|x - r_k| \leq 10^{-n}$. Next, a sequence $\{x_n\}$ of reals is said to be computable iff there is a double sequence $\{r_{kn}\}$ such that $r_{kn} \rightarrow x_n$ as $k \rightarrow \infty$ effectively in both k and n . Finally, a function of the reals is said to be computable iff it maps every computable sequence in its domain into a computable sequence and, moreover, it is effectively uniformly continuous. For details, see [Pour-el and Richards, 1989].

square. Using this embedding Moore was able to show how recursively unsolvable problems can be translated into prediction tasks about the future behavior of the particle that the Demon cannot carry out without help from an Oracle, even if it knows the initial state of the particle with absolute precision! For example, Turing's theorem says that there is no recursive algorithm to decide whether a universal Turing machine halts on a given input. Since the halting state of the universal Turing machine that has been embedded in the particle-mirror system corresponds to the particle's entering a certain region of the unit square to which it is thereafter confined, the Demon cannot predict whether the particle will ever enter this region. The generalization of Turing's theorem by Rice [1953] shows that many questions about the behavior of a universal Turing machine in the unbounded future are recursively unsolvable, and these logical questions will translate into physical questions about the behavior of the particle in the unbounded future that the Demon cannot answer without consulting an Oracle.

The reader might ask why we should fixate on the Turing notion of computability. Why not think of a deterministic mechanical system as an analogue computer, regardless of whether an abstract Turing machine can be embedded in the system? For instance, in the above example of the Newtonian particle with deterministic motion, why not say that the particle is an analogue computer whose motion "computes," for any given initial conditions $q(0), \dot{q}(0)$, the possibly non-Turing computable function $q(t) = F(q(0), \dot{q}(0), t)$? I see nothing wrong with removing the scare quotes and developing a notion of analogue computability along these lines. But the practical value of such a notion is dubious. Determining which function of t is being computed and accessing the value computed for various values of t requires ascertaining the particle position with unbounded accuracy.

Connections between non-Turing computability and general relativistic spacetimes that are inhospitable to a global version of Laplacian determinism will be mentioned below in Section 6.6.

4 DETERMINISM IN SPECIAL RELATIVISTIC PHYSICS

4.1 *How the relativistic structure of spacetime improves the fortunes of determinism*

Special relativistic theories preserve the Newtonian idea of a fixed spacetime background against which the drama of physics plays itself out, but they replace the background classical spacetimes with Minkowski spacetime. This replacement makes for a tremendous boost in the fortunes of determinism. For the symmetries of Minkowski spacetime are given by the Poincaré group, which admits a finite invariant speed c , the speed of light, making it possible to formulate laws of motion/field equations which satisfy the basic requirement that the symmetries of the spacetime are symmetries of the laws and which propagate energy-momentum no faster than c . For such laws all of the threats to classical determinism that derive from unbounded velocities are swept away.

The last point can be expounded in terms of the apparatus introduced in Section 3.8. For the type of law in question, a causal curve is a spacetime worldline whose tangent at any point lies inside or on the null cone at that point, with the upshot that domains of dependence are now non-trivial. Minkowski spacetime admits a plethora of global time functions. But in contrast with classical spacetimes, such a function t can be chosen so that the domains of dependence $D(t = \text{const})$ of the level surfaces of t are non-empty. Indeed, t can be chosen so that for each and every $t = \text{const}$ the domain of dependence $D(t = \text{const})$ is a *Cauchy surface*, i.e. $D(t = \text{const})$ is the entire spacetime. In fact, any inertial time coordinate is an example of a global time function, all of whose levels are Cauchy surfaces.⁴⁴ In the context of STR, the definition of Laplacian determinism given above in Section 2.1 is to be understood as applying to a t with this Cauchy property.

It is important to realize that these determinism friendly features just discussed are not automatic consequences of STR itself but involve additional substantive assumptions. The *stress-energy tensor* T^{ab} used in both special and general relativistic physics describes how matter-energy is distributed through spacetime. What is sometimes called the local conservation law for T^{ab} , $\nabla_a T^{ab} = 0$, where ∇_a is the covariant derivative determined by the spacetime metric, does *not* guarantee that the local energy-momentum flow as measured by any observer is always non-spacelike. That guarantee requires also that for any future pointing timelike U^a , $-T^{ab}U_a$ is a future pointing, non-spacelike vector.⁴⁵ Combining this requirement with the further demand that the local energy density as measured by any observer is non-negative, i.e. $T^{ab}U_aU_b \geq 0$ for any non-spacelike vector field U^a , produces what is called the *dominant energy condition*. Not surprisingly, this condition, together with the local conservation of T^{ab} , does guarantee that the matter fields that give rise to T^{ab} cannot travel faster than light in the sense that if T^{ab} vanishes on some spacelike region S , then it must also vanish on $D(S)$ (see [Hawking and Ellis, 1973, 91-94]). The dominant energy conditions is thought to be satisfied by all the matter-fields encountered in the actual world, but occasionally what are purported to be counterexamples appear in the physics literature.

4.2 Fundamental fields

In Section 3.4 examples were given to illustrate how fields have a hard time living up to the ideals of Laplacian determinism in classical spacetimes. The situation changes dramatically in Minkowski spacetime, which supports field equations in the form of hyperbolic pdes.⁴⁶ For example, the Klein-Gordon equation for a scalar field ϕ of mass $m \geq 0$ obeys the equation

$$(12) \quad \nabla_a \nabla^a \phi - m^2 \phi = 0$$

⁴⁴Exercise for the reader: Construct a global time function t for Minkowski spacetime such that none of the level surfaces of t are Cauchy.

⁴⁵The minus sign comes from the choice of the signature $(+ + - -)$ for the spacetime metric.

⁴⁶A standard reference on the classification of pdes relevant to physical applications is [Courant and Hilbert, 1962, Vol. 2]. See also [Beig, 2004].

which is a special case of a linear, diagonal, second order hyperbolic pde. For such equations there is a global existence and uniqueness proof for the initial value problem: given a Cauchy surface Σ of Minkowski spacetime and C^∞ initial data, consisting of the value of ϕ on Σ and the normal derivative of ϕ with respect to Σ , there exists a unique C^∞ solution of (13) throughout spacetime. Furthermore, the initial value problem is well-posed in that (in an appropriate topology) the unique solution depends continuously on the initial data. And finally the Klein-Gordon field propagates causally in that if the initial data are varied outside a closed subset $S \subset \Sigma$, the unique solution on $D(S)$ does not vary. Notice that we have a completely clean example of Laplacian determinism at work — no boundary conditions at infinity or any other enabling measures are needed to fill loopholes through which indeterminism can creep in. By contrast, giving initial data on a timelike hypersurface of Minkowski spacetime is known to lead to an improperly posed Cauchy problem; indeed, not only do solutions not depend continuously on the initial data, but there are C^∞ initial data for which there is no corresponding solution. This asymmetry between the fortunes of determinism in timelike vs. spacelike directions, could, as noted above, be turned around and used as a basis for singling out the time dimension.

It should be emphasized that only restricted classes of hyperbolic pdes are known to have well-posed initial value problems. It is a challenge to mathematical physics to show that the field equations encountered in physical theories can be massaged into a form that belongs to one of these classes. It is a comparatively easy exercise to show that, when written in terms of potentials, the source-free Maxwell equations for the electromagnetic field take the form of a system of linear, diagonal, second order hyperbolic pdes if an appropriate gauge condition is applied to the potentials. In other cases the challenge requires real ingenuity.⁴⁷

Physicists are so convinced of determinism in classical (= non-quantum) special relativistic physics that they sort “fundamental” from “non-fundamental” matter fields according as the field does or does not fulfill Laplacian determinism in the form of global existence and uniqueness theorems for the initial value problem on Minkowski spacetime. The Klein-Gordon field and the source-free Maxwell electromagnetic field qualify as fundamental by this criterion. A dust matter field, however, fails to make the cut since singularities can develop from regular initial data since, for example, in a collapsing ball of dust the density of the dust can become infinite if the outer shells fall inward fast enough that they cross the inner shells. Such shell-crossing singularities can develop even for physically reasonable initial data for the Maxwell-Lorentz equations where the source for the electromagnetic field consists of a charged dust obeying the Lorentz force law. But no great faith in determinism is needed to brush aside the failure of determinism in such examples; they can also be dismissed on the grounds that dust matter is an idealization and, like all idealizations, it ceases to work in some circumstances. Faith in determinism, however, is required to deal with what happens when the Klein-Gordon equation is converted into a non-linear equation by adding terms to

⁴⁷See [Beig, 2004] for examples.

the right hand side of (13), e.g.

$$(13) \quad \nabla_a \nabla^a \phi - m^2 \phi = \lambda \phi^2$$

where λ is a constant. It is known that solutions of (14) corresponding to regular initial data can become infinite at a finite value of t and that such data has non-zero measure (see [Keller, 1957]).

A number of attempts have been made to modify the classical Navier-Stokes equations (see Section 3.4) for dissipative fluids in order to make them consistent with STR in the sense that they become a system of hyperbolic pdes with causal propagation. A criterion of success is typically taken to be that the resulting system admits an initial value formulation, confirming once again the faith in determinism in the special relativistic setting. One difficulty in carrying out this program is that it necessitates the introduction of additional dynamical variables and additional dynamical equations, and as a result many different relativistic generalizations of the classical equations have been produced. Geroch [1995] has argued that we need not be troubled by this *embarras des riches* because the differences among the relativistic generalizations wash out at the level of the empirical observations that are captured by the Navier-Stokes theory.

4.3 Predictability in special relativistic physics

The null cone structure of Minkowski spacetime that makes possible clean examples of Laplacian determinism works against predictability for embodied observers who are not simply “given” initial data but must ferret it out for themselves by causal contact with the system whose future they are attempting to predict. Consider, for example, the predicament of an observer O whose world line is labeled γ in Fig. 1. At spacetime location p this observer decides she wants to predict what will happen to her three days hence (as measured in her proper time).




Figure 1. The failure of predictability in Minkowski spacetime

Suppose that, in fact, three-days-hence for O corresponds to the spacetime location p' . And suppose that the equations governing the variables relevant to the prediction are such that initial data on a spacelike hypersurface Σ fixes a unique solution in $D(\Sigma)$. Then to carry out the prediction by way of solving the relevant equations, O must ascertain the state on a local spacelike hypersurface that slices through the past null cone of p' , such as σ in Fig. 1. As O 's "now" creeps up her world line, the past light of the "now" sweeps out increasingly large portions of σ , but until her "now" reaches p' she does not have causal access to all of σ . And the same goes for any other local slice through the past cone of p' . Thus, the very spacetime structure that provides a secure basis for Laplacian determinism prevents O from acquiring the information she needs before the occurrence of the event that was to be predicted.

This predictability predicament can be formalized in a way that will be useful when it comes to investigating predictability in a general relativistic spacetime \mathcal{M}, g_{ab} , where \mathcal{M} is a differentiable manifold and g_{ab} is a Lorentz signature metric defined on all of \mathcal{M} , Minkowski spacetime being the special case where $\mathcal{M} = \mathbb{R}^n$ and g_{ab} is the Minkowski metric. Geroch (1977) defines the *domain of prediction* $P(q)$ of a point $q \in \mathcal{M}$ to consist of all points $p \in \mathcal{M}$ such that (i) every past directed timelike curve with future endpoint at p and no past endpoint enters the chronological past $I^-(q)$ of q ,⁴⁸ and (ii) $I^-(p) \not\subseteq I^-(q)$. Condition (i) is needed to ensure that causal processes that can influence events at p register on the region $I^-(q)$ that is causally accessible to an observer whose "now" is q , and condition (ii) is needed to ensure that from the perspective of q , the events to be predicted at p have not already occurred. The predictability predicament for Minkowski spacetime can now be stated as the theorem that for every point q of Minkowski spacetime, $P(q) = \emptyset$.

Note that the predictability predicament arises not just because of the local null cone structure of Minkowski spacetime but also because of its global topological structure. To drive home this point, suppose that space in $(1 + 1)$ -Minkowski spacetime is compactified to produce the cylindrical spacetime \mathcal{C} pictured in Fig. 2. Now predictability is possible since $I^-(q)$ for any q contains a Cauchy surface, e.g. Σ in Fig. 2. As a result $P(q) = \mathcal{C} - I^-(q)$.

For standard Minkowski spacetime and other spacetimes for which $P(q) = \emptyset$ for every spacetime point q , one can wonder how secure predictions are possible. The answer is that if complete security is required, the only secure predictions have a conditional form, where the antecedent refers to events that are not causally accessible from q . But there will be many such conditionals, with different antecedents and different consequents, and since one will not be in a position to know which of the antecedents is actualized, the best one can do is a "prediction" (all too familiar from economic forecasts) consisting of a big set of conditionals. On the other hand, if complete security is not demanded, then unconditional predictions

⁴⁸For a point q in a relativistic spacetime \mathcal{M}, g_{ab} , the chronological past $I^-(q)$ consists of all $p \in \mathcal{M}$ such that there is a future directed timelike curve from p to q . The chronological future $I^+(q)$ of a point q is defined analogously.




Figure 2. The improved fortunes of predictability when the spatial aspect of Minkowski spacetime is compactified

carrying probability instead of certainty are obtainable if inductive inference from past observations points to one of the antecedents of the set of conditionals as being highly likely.

If one wants predictions that are in principle verifiable, then a third condition needs to be added to the definition of the domain of prediction; namely, (iii) $p \in I^+(q)$. The point p in Fig. 2 satisfies clauses (i) and (ii) but not (iii).

4.4 *Special relativity and computability*

Pour-el and Richards [1981] constructed an example in which deterministic evolution does not preserve Turing computability. The equation of motion at issue is the relativistic wave equation, which in inertial coordinates is written

$$(14) \quad \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2} - \frac{\partial^2 u}{\partial t^2} = 0, \quad c \equiv 1$$

Pour-el and Richards studied solutions corresponding to initial data at $t = 0$ of the form $u(x, y, z, 0) = f(x, y, z)$, $\partial u(x, y, z, 0)/\partial t = 0$. They showed that there is a Turing computable $f(x, y, z)$ such that the corresponding solution $u(x, y, z, t)$ is not Turing computable at $t = 1$. However, such a solution is necessarily a weak solution (in the sense of Section 3.4) to the wave equation since it must be non-differentiable. And the non-preservation result is sensitive to the norm used to define convergence. Indeed, if Turing computability is defined using the

energy norm,⁴⁹ then for any Turing computable functions f and g , the solution $u(x, y, z, t)$ corresponding to $u(x, y, z, 0) = f(x, y, z)$, $\partial u(x, y, z, 0)/\partial t = g(x, y, z)$ is Turing computable (see [Pour-el and Richards, 1989, 116-118]).

5 DETERMINISM AND INDETERMINISM IN ORDINARY QM

The folklore on determinism has it that QM is the paradigm example of an indeterministic theory. Like most folklore, this bit contains elements of truth. But like most folklore it ignores important subtleties — in this instance, the fact that in some respects QM is more deterministic and more predictable than classical physics. And to distill the elements of truth from the folklore takes considerable effort — in particular, the folkloric notion that quantum indeterminism arises because the “reduction of the wave packet” is based on a controversial interpretation of the quantum measurement process. Before turning to these matters, I will discuss in Section 5.1 an issue that links to the some of the themes developed above and in Secs. 5.2-5.4 some issues unjustly neglected in the philosophical literature.

5.1 *Determinism and Galilean invariance in QM*

Here is another example of how linking determinism and symmetry considerations is fruitful in producing physical insights. Consider the motion of a single spinless particle on the real line \mathbb{R} , and work in the familiar Hilbert space \mathcal{H} of wave functions, i.e. $\mathcal{H} = L^2_{\mathbb{C}}(\mathbb{R}, dx)$. The evolution of the state $\psi(x) \in \mathcal{H}$ of the quantum particle is governed by the Schrödinger equation

$$(15) \quad i\hbar \frac{\partial \psi}{\partial t} = \hat{H}\psi$$

where \hat{H} is the Hamiltonian operator. This evolution is deterministic, or so it is said. But a puzzle is quickly generated by conjoining the presumed determinism with the presumed Galilean invariance of (16).⁵⁰ Since (16) is first order in time, giving the value of the wave function $\psi(x, 0)$ for all $x \in \mathbb{R}$ at $t = 0$ should fix the value of $\psi(x, t)$ for all $t > 0$. But how can this be if the Schrödinger equation is Galilean invariant? A proper Galilean transformation $x \rightarrow x' = x - vt$, $v \neq 0$, is the identity map for $t = 0$ but non-identity for $t > 0$. Assuming Galilean invariance of (16), this map must carry a solution to a solution. Since the map in question is the identity for $t = 0$ the two solutions should have the same initial data $\psi(x, 0)$; but since the map is not the identity for $t > 0$ the original solution and

⁴⁹For initial conditions f, g the energy norm is given by

$$\|f, g\|^2 := \int \int \int [(\nabla f)^2 + g^2] dx dy dz.$$

And for functions u on \mathbb{R}^4 the norm is $\|u(x, y, z, t)\| = \sup_t E(u, t)$, where

$$E(u, t)^2 := \int \int \int \left[\nabla u + \left(\frac{\partial u}{\partial t} \right)^2 \right] dx dy dz.$$

If u is a solution of the wave equation, then $E(u, t)$ is independent of t .

⁵⁰For a treatment of the Galilean invariance of the Schrödinger equation, see [Brown, 1999].

its image under a Galilean boost should diverge in the future, violating Laplacian determinism. The resolution of this little puzzle is to reject the implicit assumption that ψ behaves as a scalar under a Galilean transformation. In fact, Galilean invariance of the Schrödinger equation can be shown to imply that the Galilean transformation of ψ depends on the mass of the particle. And this in turn entails a “superselection rule” for mass (discovered by Bargmann [1954]) which means that a superposition of states corresponding to different masses is not physically meaningful in non-relativistic QM.

5.2 How QM can be more deterministic than classical mechanics

Physics textbooks on QM offer a procedure for quantization that starts with a Hamiltonian formulation of the classical dynamics for the system of interest and produces, modulo operator ordering ambiguities, a formal expression for the quantum Hamiltonian operator \hat{H} that is inserted into equation (16).⁵¹ But to make the formal expression into a genuine operator a domain of definition must be specified since, typically, \hat{H} is an unbounded operator and, therefore, is defined at best for a dense domain of the Hilbert space. Usually it is not too difficult to find a dense domain on which \hat{H} acts as a symmetric operator. The question then becomes whether or not this operator is essentially self-adjoint, i.e. has a unique self-adjoint (SA) extension — which will also be denoted by \hat{H} .⁵² If so, $\hat{U}(t) := \exp(-i\hat{H}t)$ is unitary for all $t \in \mathbb{R}$, and since $\hat{U}(t)$ is defined for the entire Hilbert space, the time evolve $\psi(t) := \hat{U}(t)\psi$ for every vector in the Hilbert space is defined for all times. (The Schrödinger equation (16) is just the “infinitesimal” version of this evolution equation.) Thus, if \hat{H} is essentially SA, none of the problems which beset the deterministic evolution of the classical state can trouble the deterministic evolution of the quantum state.

What is, perhaps, surprising is that the quantum Hamiltonian operator can be essentially SA in some cases where the counterpart classical system does not display deterministic evolution. Recall from Section 3.5 the example of a particle moving on the real line \mathbb{R} in a (real-valued) potential $V(x)$, $x \in \mathbb{R}$. As we saw, when the potential is proportional to $-|x|^{4/3}$ near the origin, the initial value problem for the Newtonian equation of motion does not have a unique solution. But the quantum Hamiltonian operator $\hat{H} = -\frac{\hbar^2}{2m} \frac{d^2}{dx^2} - V(x)$ is essentially SA provided that $V(x)$ is locally integrable and bounded below. And this can be satisfied by the classically non-deterministic potential by suitably modifying it away from the origin.⁵³

⁵¹See [Landsman, this vol.] for details on various approaches to quantization.

⁵²A linear operator \hat{O} defined on the dense domain $D(\hat{O})$ of the Hilbert space \mathcal{H} is *symmetric* just in case for all $\psi, \varphi \in D(\hat{O})$, $(\hat{O}\varphi, \psi) = (\varphi, \hat{O}\psi)$, where (\cdot, \cdot) is the inner product on \mathcal{H} . That \hat{O} is *self-adjoint* means that $\hat{O} = \hat{O}^*$, i.e. \hat{O} is symmetric and $D(\hat{O}) = D(\hat{O}^*)$, where \hat{O}^* is adjoint of \hat{O} . Here $D(\hat{O}^*)$ is defined to be the set of $\varphi \in \mathcal{H}$ such that there is a $\chi \in \mathcal{H}$ such that $(\hat{O}\psi, \varphi) = (\psi, \chi)$ for all $\psi \in D(\hat{O})$; then $\hat{O}^*\varphi := \chi$.

⁵³An appropriate dense domain is $\{\psi \in L^2_{\mathbb{C}}(\mathbb{R}, dx) : \psi, \psi' \in AC(\mathbb{R}) \ \& \ \hat{H}\psi \in L^2_{\mathbb{C}}(\mathbb{R}, dx)\}$ where

Another form of classical indeterminism occurs when the initial value problem has locally in time a unique solution but the solution breaks down after a finite time. The example given in Section 3.6 was that of a system of classical point mass particles moving under the influence of their mutual pairwise attractive $1/r^2$ force, and it was noted that the solution can break down either because of collision or non-collision singularities. Neither type of singularity occurs in the quantum analogue since again the quantum Hamiltonian operator for this case is essentially SA.⁵⁴

QM also cures the indeterminism of the Zeno version of classical billiards discussed in Section 3.7, at least in a backhanded sense. A putative quantum analogue would mimic the Zeno construction of an infinite number of distinct particles in the unit interval $(0, 1]$ by squeezing into that interval an infinite number of wave packets with substantial non-overlap. The latter would require that the uncertainty in position Δx associated with a wave packet becomes arbitrarily small as the origin is approached. By the uncertainty principle, the uncertainty in momentum Δp would have to become unboundedly large as the origin is approached. This represents a breakdown in the putative quantum analogue if Δx and Δp both small in comparison with some specified macroscopic standard is required for mimicking classical behavior.⁵⁵

5.3 *How QM (even without state vector reduction) can be a lot less deterministic than classical mechanics*

Determinism for the evolution of the quantum state is an all-or-nothing affair. If \hat{H} is essentially SA then the “all” alternative applies since, as already noted, the exponentiation of the unique SA extension gives a unitary evolution operator which is defined for all times and all vectors in the Hilbert space. If \hat{H} is not essentially SA there are two possibilities to consider. The first is that \hat{H} has no SA extensions. This can be dismissed on the grounds that \hat{H} should be a real operator, and every real symmetric operator has SA extensions. The second possibility is that \hat{H} has many SA extensions. Then the “nothing” alternative applies; for the exponentiations of the different SA extensions give physically distinct time evolutions. Roughly speaking, the different self-adjoint extensions correspond to different boundary conditions at the boundary points of the configuration space. Perhaps some boundary condition can be singled out and promoted to lawlike

$AC(\mathbb{R})$ stands for the absolutely continuous functions.

⁵⁴This result is known as Kato’s theorem; see [Kato, 1995, Remark 5.6]. For a more detailed discussion of the issue of essential self-adjointness and its implications for quantum determinism, see [Earman, 2005].

⁵⁵Mimicking a classical state in which a particle has given values of position and momentum requires a quantum state ψ that not only returns the given values as expectation values but also gives $(\Delta x)_\psi$ and $(\Delta p)_\psi$ small in comparison with the relevant macroscopic standard; for if $(\Delta x)_\psi$ (respectively, $(\Delta p)_\psi$) is large in comparison with the standard, there is an appreciable probability that the particle will be found with a position (respectively, momentum) different from the given value.

status, thus providing for a unique quantum dynamics. But restoring determinism by this route would require a justification for the hypothesized singling out and promotion. Alternatively, the effects of the non-essential self-adjointness of the Hamiltonian can be down played if it can be shown that the quantum dynamics associated with different self-adjoint extensions all have the same classical limit (see [Earman, 2005b]).

A toy example of the second possibility is given by a particle moving on the positive real line \mathbb{R}_+ in a potential $V(x)$, $x \in \mathbb{R}$. If the potential has the form C/x^2 , with $C > 0$, then the Newtonian initial value problem has a unique solution, and the solution is defined for all times. The intuitive explanation is that no matter how much energy it has, the classical particle cannot climb over the indefinitely high potential to reach the singularity at the origin, and it cannot reach $x = +\infty$ in finite time. However, the quantum Hamiltonian operator for this case is not essentially SA on $L^2_{\mathbb{C}}(\mathbb{R}_+, dx)$ if $C < \frac{3}{4}$ (see [Reed and Simon, 1975, Thm X.10]). The intuitive explanation is that the quantum particle can tunnel through the barrier to reach the singularity, allowing probability to leak away. This leakage is incompatible with unitary evolution, which would obtain as the result of exponentiating the unique SA extension of an essentially SA \hat{H} .

The singularity in the configuration space of the toy example is wholly artificial, being created by deleting half of the real line. But analogues in the form of naked or timelike singularities occur in general relativistic spacetimes (see Section 6.4). One can ask whether a relativistic quantum particle propagating on the background of negative mass Schwarzschild spacetime can tunnel through the effective repulsive barrier that surrounds the origin $r = 0$. Horowitz and Marolf [1995] show that the answer is positive.

Essential self-adjointness of the Hamiltonian might be promoted as a selection principle to help decide what systems are “quantum possible,” guaranteeing that (barring state vector collapse) the evolution of the quantum is deterministic. Those who think that determinism is an *a priori* truth may look favorably on this promotion, but otherwise its merits are hard to discern.

5.4 *Chaos and predictability in QM*

QM can not only be more deterministic than classical mechanics, but it can also be more predictable as well. Classical predictability is compromised or completely wrecked for observers who cannot ascertain initial conditions with complete precision if the systems whose behavior they are attempting to predict display sensitive dependence on initial conditions. But if the quantum Hamiltonian operator is essentially SA, then not only is the evolution of the quantum state completely deterministic, its predictability is not compromised by sensitive dependence on initial conditions. The point is simply that the linear, unitary nature of the evolution preserves the Hilbert space norm: $\|U(t)\psi_2 - U(t)\psi_1\| = \|U(t)(\psi_2 - \psi_1)\| = \|\psi_2 - \psi_1\|$. In words, if two states are close (as measured in the Hilbert space norm) to begin with, they remain close for all times, i.e. the evolution is stable.

This stability causes trouble for anyone seeking chaos in QM itself — they are driven to such extremes as resorting to open systems (whose evolution is not unitary) or to hidden variables whose evolution is not stable.⁵⁶ But in itself the stability of quantum evolution poses no *a priori* barrier to explaining how chaos can emerge from quantum dynamics in some appropriate classical limit. For that project only requires showing that the expectation values of relevant classical quantities can diverge rapidly enough (in some appropriate metric) to underwrite higher order ergodic properties that constitute the chaotic behavior observed on the macrolevel (see [Belot and Earman, 1997]). One obvious way to carry out this project is to use Ehrenfest’s theorem to show that in the position representation the centroid of a quantum wave packet will follow a classical trajectory, as long as the mean square width of the wave packet remains sufficiently small. However, for classically chaotic trajectories the time interval in which the latter condition holds is uncomfortably short — for example, [Zurek, 1998] estimates that for the chaotic tumbling of Hyperion (a moon of Jupiter) it is of the order of 20 years. Several authors have argued that quantum decoherence comes to the rescue (see, for example, [Zurek, 1998; 2003]), but that is a topic that is beyond the scope of this chapter. Clearly, classical chaos poses a challenge to our understanding of how the classical world emerges from quantum physics.⁵⁷ Another aspect of the classical-quantum correspondence is treated in the next section.

5.5 *State vector reduction, hidden variables, and all that*

Showing that the Hamiltonian operator \hat{H} for a quantum system of interest is essentially SA is not enough to secure the fortunes of determinism for this system, and this for two reasons. The first is that the deterministic evolution of the quantum state might be interrupted by “state vector reduction,” as is postulated in some treatments of the measurement problem in QM, by which the unitary evolution $\psi(0) \mapsto \psi(t) = \exp(-i\hat{H}t)\psi(0)$ is suspended and the quantum state jumps into an eigenstate of the observable being measured. In its crudest form state vector reduction is a literal miracle — a violation of the laws of nature — making it an inappropriate topic for the present forum. But there are more sophisticated forms of state vector reduction that model the reduction as a dynamical process. Stochastic models in which the reduction occurs discontinuously and continuously have been studied respectively by Ghirardi *et al.* [1986] and Pearle [1989]. Reduction by means of a non-linear term added to the Schrödinger equation was studied by Pearle [1976]. If the stochastic models of reduction are on the right track and if the stochastic mechanisms they postulate represent irreducible randomness, then obviously determinism is breached. By contrast, the scheme of Pearle [1976]

⁵⁶See the discussions of Kronz [1998] and Cushing and Bowman [1999]. By contrast, physicists who study “quantum chaos” do not try to find chaos in QM itself but rather study the distinguishing properties of quantum systems whose classical counterparts display chaos. For this reason Michael Berry suggested replacing “quantum chaos” with “quantum chaology.” Unfortunately the suggestion did not stick.

⁵⁷For a comprehensive survey of this problem, see [Landsman, this vol.].

achieves a deterministic state vector reduction with the help of hidden variables.⁵⁸ None of these alternatives to standard non-relativistic QM has been generalized to a viable relativistic quantum field theory, and as far as I am aware none of them play any role in the main lines of research on quantum gravity that come from string theory or loop quantum gravity (see Section 8). Thus, at present it does not seem productive to speculate about the implications for determinism of possible modifications to QM that may or may not become part of some future physics. However, the motivation for introducing state vector reduction is relevant here, for it leads to the second set of reasons why the conventional quantum state dynamics may not be sufficient to secure determinism for the quantum domain.

Classical (= non-quantum) theories wear their interpretations on their sleeves.⁵⁹ For example, for a classical theory that admits an (unconstrained) Hamiltonian formulation, observables are in one-one correspondence with functions from the phase space $\Gamma(q, p)$ to (say) the reals \mathbb{R} . The intended interpretation is that if f_O is the function corresponding to the observable O , then the value of O at t is o iff the state $(q(t), p(t))$ at t is such that $f_O(q(t), p(t)) = o$. This scheme can be liberalized to allow for dispositional observables which have definite values in only some states; for such an O the representing function f_O is only a partial function. Another liberalization is to allow that the range of f_O includes “fuzzy” (e.g. interval) values. To get an interpretation of QM along similar lines requires adding to the formalism at least two things: (i) an account of which SA operators correspond to quantum observables, and (ii) a semantics for quantum observables in the form of a value assignment rule that specifies what values the observables take under what conditions. I will simply assume that part (i) of the interpretation problem has been solved.

The most obvious way to supply part (ii) would be to ape the classical value assignment rule, replacing the classical state space $\Gamma(q, p)$ by the quantum state space to get a value assignment rule of the form: the value of quantum observable O at t is o iff the state vector $\psi(t)$ is such that $f_O(\psi(t)) = o$ where f_O is the representing function for the quantum observable O . If, as implicitly assumed in this formulation, the quantum state space is taken to be the the unit sphere $S\mathcal{H}$ of the Hilbert space \mathcal{H} of the system (i.e. $\{\psi \in \mathcal{H} : (\psi, \psi) = 1\}$), then as far as standard QM is concerned, gauge freedom is present since any two elements of $S\mathcal{H}$ that differ by a phase factor correspond to the same physical state in that all expectation values of observables are the same for the two quantum states. This gauge redundancy can be removed by taking the state space to be the projective Hilbert space $\mathbb{P}\mathcal{H}$, defined as the quotient of $S\mathcal{H}$ by the action of $\psi \mapsto \varsigma\psi$ where $\varsigma \in \mathbb{C}$ with $|\varsigma| = 1$; equivalently, $\mathbb{P}\mathcal{H}$ is the space of rays or one-dimensional subspaces of \mathcal{H} . Thus, from the point of view of conventional QM, the value assignment rule should obey the restriction that $f_O(\psi) = f_O(\psi')$ whenever the

⁵⁸The hidden variables are the phase angles, an idea revived by Ax and Kochen [1999]; see below.

⁵⁹But recall from Section 3 that if determinism is demanded, then the initial on-the-sleeve interpretation may have to be modified by seeing gauge freedom at work.

unit vectors ψ' and ψ belong to the same ray. Allowing the value assignment to depend on the phase would amount to introducing “hidden variables,” in the terminology used below.

In any case, if the quantum value assignment rule takes the form under discussion and if the problems discussed in Section 5.3 are waived, then arguably QM is a deterministic theory, and this is so even if f_O is a partial function (i.e. is undefined for some quantum states) or even if f_O can take fuzzy values. For assuming no state vector reduction, the state $\psi(0)$ at $t = 0$ uniquely determines the state $\psi(t)$ at any $t > 0$; and assuming the implementation of (ii) under discussion, the state $\psi(0)$ at $t = 0$ uniquely determines the value assignments at any later time $t > 0$. That an observable is assigned no value or a fuzzy value at t does not represent a failure of determinism, which requires only that the laws plus the initial state determine the present and future values of all observables *to the extent that these values are determinate at all*. Thus, on the present option it is a mistake to view the Kochen-Specker theorem, and other subsequent no-go theorems, as showing that QM does not admit a deterministic interpretation. Rather, what these no-go results show is that, subject to certain natural constraints,⁶⁰ some subset of quantum observables cannot all be assigned simultaneously sharp values.⁶¹ The same goes for the Bell-type theorems, which are best interpreted as extensions of the Kochen-Specker no-go result to an even smaller set of observables (see [Fine, 1982a; 1982b]).

One example of the type of value assignment rule at issue is the eigenvalue-eigenvector rule which says that, for an observable O whose quantum operator \hat{O} has a discrete spectrum, O has a sharp value at t iff $\hat{O}\psi(t) = o\psi(t)$, in which case $O(t) = o$. But it is just this eigenvalue-eigenvector link that leads to the notorious measurement problem in QM in the form of the inability of the theory to explain why measurements have definite outcomes, and it is this problem that motivated the idea of state vector reduction. In essence the problem arises because of the insistence that “measurement” should not be taken as a primitive term but should be analyzed within QM itself as a physical interaction between the object system and a measuring instrument. But while the application of the standard linear, unitary dynamics to the composite object-system + measurement-apparatus-system can establish a one-one correlation between the eigenstates of the object observable of interest and the eigenstates of the “pointer observable” of the measuring instrument, the application of the eigenvector-eigenvalue rule to the post measurement composite state function yields the unacceptable result that the “pointer” on the measuring instrument is not pointing to any definite value

⁶⁰For example, it is natural to require that if the quantum value assignment rule for O assigns O a definite value, that value lies in the spectrum of the operator \hat{O} corresponding to O . And it is natural to require that for suitable functions g , $F_{g(O)} = g \circ F_O$.

⁶¹It follows from Gleason’s theorem that, subject to very plausible constraints on value assignments, not all of the (uncountably infinite number of) self-adjoint operators in a Hilbert space of dimension 3 or greater can be assigned simultaneously definite values belonging to the spectra of these operators. The Kochen-Specker theorem shows that the same conclusion can be drawn for a finite set of quantum observables. See [Redhead, 1987] for an account of these no-go results.

(see [Albert, 1992] for a detailed exposition). The Schrödinger cat paradox is a cruel illustration of this conundrum in which “live cat” and “dead cat” serve as the “pointer positions.”

Thus, if the eigenvalue-eigenvector link is maintained, there are compelling reasons to contemplate a modification of the standard quantum dynamics in order to ensure that in measurement interactions the quantum state evolves into an eigenstate of the relevant observables. But since the decision was made above not to treat such modifications, the discussion that follows will be confined to the other option, namely, the use of a value assignment rule that breaks the eigenvalue-eigenvector link, possibly with the help of “hidden variables” that supplement the quantum state. If hidden variables X are used, the value assignment rule takes the form: the value of quantum observable O at t is o iff the total state $(\psi(t), X(t))$ is such that $f_O(\psi(t), X(t)) = o$, where again f_O stands for the representing function for the observable O but is now a function defined on the augmented state space. If the evolution of the total state is deterministic, then by the same argument as before, the quantum domain is fully deterministic if QM is true. An example is supplied by the Bohm interpretation where $X(t)$ specifies the positions of the particles at t . The quantum component of the total state evolves according to Schrödinger dynamics, and the postulated equation of motion for particle positions guarantees that $(\psi(0), X(0))$ uniquely determines $(\psi(t), X(t))$ for $t > 0$ (see [Holland, 1993] for a survey). On the Bohm interpretation many quantum observables have a dispositional character, taking on determinate values only in adequately specified contexts (typically including a measurement apparatus together with its hidden variables). For example, in the context of a Stern-Gerlach experiment a spin $\frac{1}{2}$ particle will have spin-up (or spin down) just in case the position of the particle lies the appropriate region of the apparatus. The validity of the claim that the Bohm interpretation resolves the measurement problem thus turns on whether all measurements can be reduced to position measurements.

The family of modal interpretations of QM also attempt to resolve the measurement problem by breaking open the eigenvalue-eigenvector link wide enough to allow measurements to have definite outcomes but not so wide as to run afoul of the Kochen-Specker type impossibility results (see [Bub, 1997] and [Dickson, this vol.] for overviews), but in contrast to the Bohm interpretation the modal interpretations have no commitment to maintaining determinism. Very roughly the idea is that an observable associated with a subsystem of a composite system in state $\psi(t)$ has a definite value just in case the reduced density matrix of the subsystem is a weighted sum of projectors associated with an eigenbasis of the observable. This guarantees that in an idealized non-disturbing measurement interaction in which the pointer positions of the measuring instrument are perfectly correlated with the possible values of the object system observable being measured, both the pointer observable and the object system observable have definite values.⁶²

⁶²More generally, the interaction of a system with its environment will mean that “measurement” of the system is going on all the time. Thus, decoherence aids the modal interpretation by providing the conditions of applicability of the interpretation. In the other direction, deco-

Most forms of the modal interpretation supply the probabilities for an observable to have particular values, assuming that the conditions are appropriate for the observable to have a determinate value; but they are silent as to what the actual value is. Nevertheless, the actual possessed values of quantum observables can be taken to play the role of the hidden variables X , and one can ask whether the total state (ψ, X) can be given a deterministic dynamics. The answer is negative for versions of the modal interpretation discussed in the philosophy literature since these versions do not supply enough hidden variables to allow for determinism. For example, at the time $t > 0$ when an ideal measurement interaction is completed and the eigenstates of pointer position are perfectly correlated with eigenstates of the object observable, the standard modal interpretations say that both the object observable and the pointer observable have definite values. In different runs of the experiment these correlated observables have different values. But in all the runs the initial quantum state $\psi(0)$ is the same, and the experimental situation can be arranged so that modal interpretations say that the initial possessed values $X(0)$ are the same. This failure of determinism is of no concern to the modal theorist whose goal is to solve the measurement problem. To this end it is enough to show that there is a stochastic dynamics for possessed values that is compatible with the statistical predictions of QM. In fact, there is a vast array of such dynamics (see [Dickson, 1997] and [Bacciagaluppi, and Dickson, 1998]).

A different version of the modal interpretation, proposed by Ax and Kochen [1999], takes the option mentioned above of extending the standard quantum state space of rays $\mathbb{P}\mathcal{H}$ to unit vectors $\mathcal{S}\mathcal{H}$. Elements of the former are supposed to characterize statistical ensembles of systems while elements of the latter characterize individual systems. This extension allows the modal interpretation to specify what value an observable has, in circumstances when it has a definite value, and also to provide for a deterministic evolution of the augmented quantum state. It is postulated that the ensemble corresponding to a ray $\varsigma\psi$, $|\varsigma| = 1$, is composed of individual systems with phase factors ς having an initial uniformly random distribution, which accounts for the apparent failure of determinism.

Both the Bohm interpretation and the family of modal interpretations have difficulties coping with relativistic considerations. The former does not have any natural generalization to QFT, at least not one which takes seriously the lesson that in QFT fields are the fundamental entities and particles are epiphenomena of field behavior. The latter does possess a natural generalization to QFT, but it yields the unfortunate consequence that in situations that are standardly discussed, no subsystem observable has a definite value (see [Earman and Ruetsche, 2005] and the references therein).

Many worlds interpretations of QM can be given a literal or a figurative reading (see [Barrett, 1999] for an overview). On the literal reading there are literally many worlds in that spacetime splits into many branches which, from the branch

herence requires something akin to the modal interpretation, for otherwise it does not, contrary to the claims of its promoters, resolve the measurement problem. For more on decoherence, see [Landsman, this vol.].

time onwards, are topologically disconnected from one another (see, for example, [McCall, 1995]).⁶³ This form of many worlds can be described as a hidden variable interpretation by taking the hidden variables X to describe the spacetime branching and by taking the representing function f_O to be a mapping from the total state $(\psi(t), X(t))$ to a vector, possibly with infinitely many components labeled α , where the component α supplies the value at t of O in branch α . The fate of determinism then depends on whether or not the story of when and how branching takes place makes the evolution of the total state (ψ, X) deterministic. On the figurative reading of “many worlds” there is literally only one world, but there are many minds, contexts, perspectives, or whatever. Also there is no such thing as an observable O simpliciter but rather an observable O -in-context- α , denoted by O_α . If the representing function f_{O_α} is a function of the quantum state only, then determinism seems to be secured. However, our notation is defective in disguising the need for a specification of the contexts that are available at any given time. That specification is determined by the quantum state $\psi(t)$ alone if there is a “democracy of bases,” i.e. any “branch” of $\psi(t)$ expressed as a linear combination of the vectors of any orthonormal basis of the Hilbert space of the system defines a context. Such a radical democracy seems incompatible with experience, e.g. in the Schrödinger cat experiment we either see a live cat or we see a dead cat, and we never experience a superposition of seeing a live and seeing a dead cat.⁶⁴ To overcome this difficulty some many world theorists propose to work with a preferred set of bases. The issue of determinism then devolves on the question of whether the specification of the set of preferred bases is deterministic. Even if the many worlds interpretation — on either the literal or figurative version — secures ontological determinism, the price seems to be a radical epistemic indeterminism: How do I know which branch of a splitting world or which context of a non-splitting world I am in? Being told that there is no “I” only an “I-in-branch- α ” or an “I-in-context- α ” is of no help when I — whichever I that is — have to make a prediction about the outcome of a measurement. Here all I can do is fall back on the statistical algorithm of QM. The many worlds interpretation seems to guarantee that even if the world is ontologically deterministic, it behaves, as far as anyone can judge, as if there is an irreducible stochasticity.

Although the discussion of the quantum measurement problem and its ramifications has been very sketchy, I trust it is sufficient to indicate why it is vain to hope for a simple and clean answer to the question of whether the truth of QM entails the falsity of determinism. To arrive at an answer to that question calls for winnowing the various competing interpretations of QM, a task that is far from straightforward, especially since judgments about how to perform the winnowing

⁶³How to describe branching spacetimes within the context of conventional spacetime theories is a ticklish matter. Perhaps the most promising move is to hold on to the assumption that spacetime is a differentiable manifold but abandon the assumption that it is a Hausdorff manifold. However, non-Hausdorff manifolds can display various pathological properties that threaten determinism, e.g. geodesics can bifurcate. See Section 6.1.

⁶⁴But how can we be sure? Perhaps momentary mental confusion is a superposition phenomenon.

are inevitably colored by attitudes towards determinism.

6 DETERMINISM IN CLASSICAL GTR

6.1 *Einstein's revolution*

Einstein's GTR was revolutionary in many respects, but for present purposes the initially most important innovation is that GTR sweeps away the notion — shared by all pre-GTR theories — of a fixed spacetime background against which the evolution of particles and fields takes place. In GTR the spacetime metric is a dynamical field whose evolution is governed by Einstein's gravitational field equations (EFE). Before discussing the issue of whether this evolution is deterministic, two preliminary matters need attention.

First, general relativists typically assume that the manifold \mathcal{M} of a relativistic spacetime \mathcal{M}, g_{ab} is Hausdorff.⁶⁵ Without this stipulation determinism would be in deep trouble. For example, non-Hausdorff spacetimes can admit a bifurcating geodesics; that is, there can be smooth mappings γ_1 and γ_2 from, say, $[0, 1]$ into \mathcal{M} such that the image curves $\gamma_1[0, 1]$ and $\gamma_2[0, 1]$ are geodesics that agree for $[0, b)$, $0 < b < 1$, but have different endpoints $\gamma_1(1)$ and $\gamma_2(1)$. According to GTR, the worldline of a massive test particle not acted upon by non-gravitational forces is a timelike geodesic. But how would such a particle know which branch of bifurcating geodesic to follow? Additionally, the (local) uniqueness of solutions to the initial value problem for EFE discussed below in Section 6.3 would fail if non-Hausdorff attachments were allowed.

Second, the reader is reminded that attention is being restricted to relativistic spacetimes \mathcal{M}, g_{ab} that are temporally orientable, and it is assumed that one of the orientations has been singled out as giving a directionality to time. But even with this restriction in place, some of the spacetimes embodied in solutions to EFE are inimical to the formulation of global Laplacian determinism given in Section 2.1. For example, such spacetimes may not admit a global time function. Indeed, the spacetime of the Gödel cosmological model not only does not admit a global time function, but it does not even admit a single global time slice (spacelike hypersurface without edges) so that one cannot meaningfully speak of the universe-at-a-given-moment.⁶⁶

One response would be to narrow down the class of physically acceptable models of GTR by requiring that, in addition to satisfying EFE, such models must also fulfill restrictions on the global causal structure of spacetime that rule out such monstrosities as Gödel's model and other models which contains closed timelike

⁶⁵ \mathcal{M} is *Hausdorff* iff for any $p, q \in \mathcal{M}$ with $p \neq q$, there are neighborhoods $N(p)$ and $N(q)$ such that $N(p) \cap N(q) = \emptyset$. Of course, a manifold is (by definition) locally Euclidean and, therefore, locally Hausdorff.

⁶⁶This is a consequence of three features of Gödel spacetime: it is temporally orientable (i.e. it admits a continuous non-vanishing timelike vector field), it is simply connected, and through every spacetime point there passes a closed future directed timelike curve. For a description of the Gödel solution, see [Hawking and Ellis, 1973, 168–170] and [Malament, 1984].

curves. This move has the independent motivation of avoiding the “paradoxes of time travel.”⁶⁷ But much stronger causality conditions are needed to underwrite the global version of Laplacian determinism in the general relativistic setting.

In the first place, to carry over the classical conception of Laplacian determinism to the context of a general relativistic spacetime requires that the spacetime admit a global time function, which is not guaranteed by the absence of closed timelike curves. But even the requirement of a global time function is not strong enough because it provides no guarantee that the level surfaces of *any* such function will have the Cauchy property. To be at home, Laplacian determinism requires a spacetime \mathcal{M}, g_{ab} that is *globally hyperbolic*, which is the conjunction of two conditions: first, \mathcal{M}, g_{ab} must be *strongly causal* in that for any $p \in \mathcal{M}$ and any neighborhood p there is a subneighborhood such that once a future directed causal curve leaves, it never reenters (intuitively, there are no almost closed causal curves); and second, for every $p, q \in \mathcal{M}$, the causal diamond $J^+(p) \cap J^-(q)$ is compact.⁶⁸ Global hyperbolicity guarantees that \mathcal{M}, g_{ab} can be foliated by Cauchy surfaces and that \mathcal{M} is diffeomorphically $\Sigma \times \mathbb{R}$, where Σ is an $n - 1$ dimensional manifold if $\dim(\mathcal{M}) = n$. But simply stipulating global hyperbolicity has all the virtues of theft over honest toil. So let us see what can be achieved by honest toil.

6.2 Determinism and gauge freedom in GTR

For pre-relativistic theories a constant theme was that creating an environment friendly to determinism requires willingness to either beef up the structure of the background spacetime or else to see gauge freedom at work in sopping up the apparent indeterminism (recall Section 3.3). But in GTR there is no fixed background structure. Thus, one would expect that GTR either produces indeterminism or else that there is a non-trivial gauge symmetry at work. This expectation is not disappointed.

To see why it is necessary to be more detailed about the EFE:

$$(16) \quad R_{ab} - \frac{1}{2}Rg_{ab} + \Lambda g_{ab} = \kappa T_{ab}$$

where R_{ab} and $R := R^c_c$ are respectively the Ricci tensor (which is defined in terms of g_{ab} and its derivatives) and the Ricci scalar, Λ is the cosmological constant, and T_{ab} is the stress-energy tensor. The cosmological constant can be ignored for present purposes, but it is currently the object of intense interest in cosmology since a positive Λ is one of the candidates for the “dark energy” which is driving the accelerating expansion of the universe (see [Ellis, this vol.]).

A potential model of the theory is then a triple $\langle \mathcal{M}, g_{ab}, T_{ab} \rangle$ where g_{ab}, T_{ab} satisfy (17) at all points of \mathcal{M} . Building such a model seems all too easy: start

⁶⁷But see [Earman, Smeenk, and Wüthrich, 2005] which argues that the so-called paradoxes of time travel do not show that time travel is conceptually or physically impossible.

⁶⁸ $J^+(p)$ (respectively, $J^-(p)$) denotes the *causal future* (respectively, *causal past*) of p , i.e., the set of all points q such that there is a future directed causal curve from p to q (respectively, from q to p).

with any any general relativistic spacetime \mathcal{M}, g_{ab} , compute the Einstein tensor $G_{ab} := R_{ab} - \frac{1}{2}Rg_{ab}$, and define the stress-energy tensor by $T_{ab} := \kappa G_{ab}$. Thus, the understanding must be that T_{ab} arises from a known matter field. And in order to make the environment as friendly as possible for determinism, it will be assumed that the T_{ab} 's that are plugged into the right hand side of (17) fulfill the dominant energy condition (see Section 4.1) which, together with the local conservation law $\nabla^a T_{ab} = 0$ (which itself is a logical consequence of (17)), guarantees that matter-energy does not propagate faster than light.

Even with these enabling stipulations in place, it seems at first glance that determinism gets no traction, at least not if a naively realistic interpretation is given to the models of the theory. The difficulty can be made apparent by repeating a variant of the construction given in Section 3.2. Let $\langle \mathcal{M}, g_{ab}, T_{ab} \rangle$ be a model satisfying all of the above stipulations, and suppose that the spacetime \mathcal{M}, g_{ab} satisfies all of the causality conditions you like, e.g. that it is globally hyperbolic. Since there is no fixed background structure to respect, save for the topological and differentiable structure of \mathcal{M} , one is free to choose a diffeomorphism $d : \mathcal{M} \rightarrow \mathcal{M}$ such that d is the identity map on and to the past of some Cauchy surface Σ of \mathcal{M}, g_{ab} but non-identity to the future of Σ . Then $\langle \mathcal{M}, d^*g_{ab}, d^*T_{ab} \rangle$, where d^* indicates the drag along by d , will also be a model satisfying all of the stipulations imposed on $\langle \mathcal{M}, g_{ab}, T_{ab} \rangle$. By construction, $d^*g_{ab}(p) = g_{ab}(p)$ and $d^*T_{ab}(p) = T_{ab}(p)$ for all p on or to the past of Σ , but $d^*g_{ab}(p) \neq g_{ab}(p)$ and $d^*T_{ab}(p) \neq T_{ab}(p)$ for some points p to the future of Σ (unless we have inadvertently chosen a d that is a symmetry of g_{ab} and T_{ab} , which can always be avoided). The existence of this pair of models that agree for all past times but disagree in the future is a violation of even the weakest cousin of Laplacian determinism, at least if the surface structure of the theory — tensor fields on a manifold — is taken at face value.

When this sort of construction threatened to undermine determinism in a pre-GTR setting, two options were available for shoring up determinism: add more structure to the background spacetime or reject the container view of spacetime. The first option is ineffective unless the additional elements of spacetime structure are non-dynamical objects, but this represents a retreat from one of the key features Einstein's revolution. If there is to be no retreat, then the second option must be exercised. In the present context the option of rejecting the container view of spacetime takes the form of rejecting the naive realism that reads the theory as describing tensor fields living on a manifold.

Choosing the second option has a principled motivation which is not invented to save determinism in GTR but which follows in line with the treatment of gauge symmetries in pre-general relativistic theories. The field equations (17) of GTR are the Euler-Lagrange equations derived from an action principle that admits the diffeomorphism group as a variational symmetry group. Thus, Noether's second theorem applies, indicating that we have a case of underdetermination — more "unknowns" than there are independent field equations — and that arbitrary functions of the spacetime variables will show up in solutions to the field equations.

Switching from the Lagrangian to the Hamiltonian formulation, it is found, as expected, that GTR is a constrained Hamiltonian theory. There are two families of first class constraints, the momentum constraints and the Hamiltonian constraints.⁶⁹ Unfortunately the Poisson bracket algebra of these constraints is not a Lie algebra,⁷⁰ and consequently one cannot directly associate the diffeomorphism group, which acts on the spacetime, with a group which acts on the Hamiltonian phase space by finding a natural homomorphism of the Lie algebra of the diffeomorphism group into the constraint algebra. This glitch is overcome by Isham and Kuchař [1986a; 1986b] who show that if appropriate embedding variables and their conjugate momenta are used to enlarge the phase space, then the enlarged constraint algebra is a Lie algebra, and that there exists a homomorphism of the Lie algebra of the spacetime diffeomorphism group into the new constraint algebra. Thus, the standard apparatus for treating gauge symmetries can be applied, yielding the verdict that the diffeomorphism invariance of GTR is to be interpreted as a gauge symmetry. On this interpretation, the above construction does not demonstrate that GTR is indeterministic but rather produces a *faux* violation of determinism by taking advantage of the redundancy of the surface structure theory in the sense of the many-to-one correspondence between the surface structure models and the intrinsic physical situation they describe; in particular, the models $\langle \mathcal{M}, d^*g_{ab}, d^*T_{ab} \rangle$ and $\langle \mathcal{M}, g_{ab}, T_{ab} \rangle$ in the above construction cannot threaten determinism since they are to be interpreted as different descriptions of the same physical situation. Of course, the apparatus at issue has built into it a commitment to determinism, so its application to GTR cannot be taken as part of a proof that the correct interpretation of GTR makes it a deterministic theory. The only claim being made here is that this determinism-saving move for GTR is not ad hoc but is part of a systematic approach to gauge symmetries that is taken to yield the “correct” results for pre-GTR theories.⁷¹

What is so clear using hindsight wisdom took Einstein many years of struggle to understand. His infamous “hole argument” can be seen as a discovery of this underdetermination problem.⁷² What muddied the waters was a confusion between two senses of general covariance. *Formal general covariance* demands that the laws of motion/field equations of a theory be written in a form that makes them

⁶⁹The plural is used here since there is a momentum constraint and a Hamiltonian constraint for every point of space.

⁷⁰The bracket of a pair of the constraints is not always a linear combination of the constraints multiplied by a “structure constant.” This failure of the constraint algebra to form a Lie algebra means that GTR is not a gauge theory in the sense of Yang-Mills. But it certainly does not mean that GTR does not contain non-trivial degrees of gauge freedom.

⁷¹Because these matters are surrounded by so much controversy in the philosophical literature, I want to emphasize as strongly as possible that I am not proposing a new way of looking at GTR but am simply expounding what is the standard view among general relativists; see, for example, [Wald, 1984].

⁷²See [Norton, 1984] and [Stachel, 1986] for accounts of how the “hole argument” figured in Einstein’s search for his gravitational field equations. And see [Rovelli, this. vol, Ch. 12] for an account of how reflecting on the lessons of the “hole argument” influenced his understanding of classical GTR and his approach to quantum gravity.

covariant under arbitrary coordinate transformations. The terminology “formal” was chosen with malice aforethought since the demand of formal general covariance is a demand on the form rather than on the content of theory. For example, Newtonian and special relativistic theories can be reformulated, without change of content, so as to meet this demand. Indeed, the fact that Newtonian and special relativistic theories can be formulated in a completely coordinate-free manner already should make it clear that coordinates cannot matter.⁷³ *Substantive general covariance* demands diffeomorphism invariance (e.g. that for arbitrary diffeomorphism of \mathcal{M} , $(\mathcal{M}, d^*g_{ab}, d^*T_{ab})$ is a model of the theory if $(\mathcal{M}, g_{ab}, T_{ab})$ is) and that this diffeomorphism invariance is a gauge symmetry. Again the terminology “substantive” was chosen with malice aforethought since the demand of substantive general covariance is not automatically met, without change of content, for formally generally covariant Newtonian and special relativistic theories, at least not by the lights of apparatus for treating gauge symmetries that has been touted here (see [Earman, 2006]).

What invites confusion is the fact that a spacetime coordinate transformation can be taken to indicate either a relabeling of spacetime points or as indicating a (local) diffeomorphism. In the first guise these transformations are gauge transformations of an innocent kind: they relate the various coordinate representations of the intrinsic coordinate-free objects g_{ab} and T_{ab} obtained by taking the components of these objects in different coordinate systems. But there is nothing new here as regards GTR since exactly the same story holds for intrinsic coordinate-free presentations of pre-GTR theories. In the second guise, however, these transformations may or may not be gauge transformations — it depends on the content of the theory.

When he first discovered the underdetermination problem by means of the “hole argument,” Einstein took it to reveal a real and intolerable form of underdetermination. To avoid it, he thought he had to reject formal general covariance as a *desideratum* for gravitational field equations. Only after wandering in the wilderness of non-covariant equations for almost three years did he re-embrace general covariance. In effecting the re-embrace Einstein did not speak the language of gauge symmetries (the terminology and the apparatus had not been invented), so he did not say that the gauge interpretation of GTR lowers the hurdle for determinism in that it requires only the uniqueness of evolution for gauge invariant quantities. But he said what amounts to the same thing; or rather he said it for a subclass of the gauge invariant quantities of GTR — what are called “point coincidences,” i.e. things like the intersection of light rays.⁷⁴

Many philosophers have traced Einstein’s path in various ways. Very few of them, however, have owned up to the implications of where the path leads. If

⁷³In the above presentation I have intentionally used the “abstract index” notation. Thus, g_{ab} stands for a symmetric, covariant tensor field that is defined in a coordinate-free manner as a bilinear map of pairs of tangent vectors to \mathbb{R} . This object can be represented by its coordinate components g_{jk} in a coordinate system $\{x^i\}$. The transformations between two such representations are gauge transformations, albeit trivial ones.

⁷⁴See [Howard, 1999] for an account of Einstein’s use of this term.

determinism in GTR is saved by treating diffeomorphism invariance as a gauge symmetry, then the only “observables” (= genuine physical magnitudes) of GTR are gauge invariant quantities. This is easy enough to say, but what exactly is the nature of the gauge invariant structure that underlies the surface structure? This is a crucial issue for those physicists who pursue a quantum theory of gravity by applying some version of the canonical quantization program to GTR, for on this program it is the “observables” of classical GTR that will be turned into quantum observables in the sense of self-adjoint operators on the Hilbert space of quantum gravity. There is no standard answer to the question of how best to characterize the observables of classical GTR. But one thing is sure: none of the familiar quantities used in textbook presentations of GTR, not even scalar curvature invariants such as the Ricci scalar appearing in the field equations (17), count as observables in the sense under discussion. And more particularly, no local quantities — quantities attached to spacetime points or finite regions — are gauge invariants. In this respect the gauge-free content of the theory has a non-substantialist flavor. Whether this content can be characterized in a way that also satisfies traditional relationalist scruples remains to be seen.

A second closely related implication of treating the diffeomorphism invariance of GTR as a gauge symmetry concerns the nature of time and change. In the Hamiltonian formalism the implication takes the form of a “frozen dynamics.” Applying to the Hamiltonian constraint of GTR the doctrine that first class constraints generate gauge transformations leads directly to the conclusion that motion in GTR is pure gauge. Put another way, the instantaneous states in the Hamiltonian formulation of the theory contain redundant structure, and any two such states, where one is generated from another by solving the Hamiltonian form of EFE, are equivalent descriptions of the same intrinsic, gauge invariant situation.⁷⁵

For those who find these implications unpalatable, the heterodox moves mentioned in Section 3.3 may be attractive. As far as I am aware, however, such heterodoxy as applied to GTR has not been seriously pursued by the physics community.

6.3 *The initial value problem in GTR*

For the sake of simplicity consider the initial value problem for the source-free or vacuum EFE, i.e. (17) with $T_{ab} \equiv 0$. Since these equations are second order in time, presumably the appropriate initial data consist of the values, at some given time, of the spacetime metric and its first time derivative. The technical formulation of this idea is to take an initial data set to consist of a triple (Σ, h_{ab}, k_{ab}) , with the following features and intended interpretations. Σ is a three-manifold, which is to be embedded as a spacelike hypersurface of spacetime \mathcal{M}, g_{ab} . h_{ab} is a smooth Riemann metric on Σ , which will coincide with the metric induced on Σ by the spacetime metric g_{ab} when Σ is embedded as a spacelike hypersurface

⁷⁵For more on the problem of time in GTR and quantum gravity, see [Belot and Earman, 1999], [Belot, this vol.], and [Rovelli, this vol.].

of \mathcal{M}, g_{ab} . And k_{ab} is a smooth symmetric tensor field on Σ that coincides with the normal derivative of h_{ab} when Σ is embedded as a spacelike hypersurface of \mathcal{M}, g_{ab} . A spacetime \mathcal{M}, g_{ab} that fulfills all of these roles is said to be a *development* of the initial data set (Σ, h_{ab}, k_{ab}) . If the development \mathcal{M}, g_{ab} of the initial data set (Σ, h_{ab}, k_{ab}) satisfies the source-free EFE, then h_{ab} and k_{ab} cannot be specified arbitrarily but must satisfy a set of constraint equations. The existence and uniqueness result for the source-free EFE takes the following form⁷⁶: Let (Σ, h_{ab}, k_{ab}) be an initial value set satisfying the constraint equations; then there exists a development \mathcal{M}, g_{ab} of the initial data that is the unique — up to diffeomorphism — maximal Cauchy development satisfying the source-free field equations. Furthermore, g_{ab} depends continuously on the initial data (see [Hawking and Ellis, 1973] for details of the relevant topology).

Just as the proof of the well-posedness of the initial value problem for the homogeneous Maxwell equations exploits the gauge freedom in the electromagnetic potentials (see Section 4.2), so the existence and uniqueness proof for EFE exploits the idea that diffeomorphism invariance is a gauge symmetry of GTR. When the metric potentials g_{ij} (i.e. the coordinate components of the metric g_{ab}) are subjected to a gauge condition (called the harmonic coordinate condition), the EFE take the form of a system of quasi-linear, diagonal, second order hyperbolic pdes, which are known to have locally well-posed initial value problem.

That the development \mathcal{M}, g_{ab} of the given initial data is a Cauchy development means that Σ is a Cauchy surface of \mathcal{M}, g_{ab} (and, thus, that this spacetime is globally hyperbolic). That it is the maximal Cauchy development means that there is no proper extension of \mathcal{M}, g_{ab} which is a solution of the source-free EFE and for which Σ is a Cauchy surface. The up-to-diffeomorphism qualifier to uniqueness was to be expected from the discussion of gauge freedom in the previous subsection, and in turn the presence of this qualifier shows that the heuristic discussion given there can be given precise content. Here the qualifier means that if \mathcal{M}', g'_{ab} is any other maximal development satisfying the source-free EFE, then there is a diffeomorphism $d : \mathcal{M} \rightarrow \mathcal{M}'$ such that $d^* g_{ab} = g'_{ab}$.

Curie's Principle (see Section 2.3 above) would lead one to believe that a symmetry of the initial value set (Σ, h_{ab}, k_{ab}) for the vacuum EFE should be inherited by the corresponding solution. And so it is. Let $\varphi : \Sigma \rightarrow \Sigma$ be a diffeomorphism that is a symmetry of the initial data in the sense that $\varphi^* h_{ab} = h_{ab}$ and $\varphi^* k_{ab} = k_{ab}$. Then as shown by Friedrich and Rendall [2000, 216–217], if ψ is an embedding of Σ into the maximal Cauchy development determined by (Σ, h_{ab}, k_{ab}) , there exists an isometry $\bar{\psi}$ of this development onto itself such that $\bar{\psi} \circ \varphi = \varphi \circ \psi$, i.e. there is an isometry of the maximal Cauchy development whose restriction to $\varphi(\Sigma)$ is ψ . Moreover, this extension of the symmetry of the initial data is unique.

The initial value problem for the sourced EFE involves not only the stress-energy tensor T_{ab} but also the equations of motion for the matter fields that give rise to T_{ab} and, in particular, the coupling of these matter fields to gravity and to each other. Whether the coupled Einstein-matter equations admit an initial value formulation

⁷⁶This formulation is taken from Wald [1984, Theorem 10.2.2].

and, if so, whether the initial value problem is well-posed are issues that have to be studied on a case-by-case basis. For what seem to be appropriate choices of coupling, the initial value problem for the combined Einstein-Klein-Gordon equations and the Einstein-Maxwell equations have existence and uniqueness results similar to that for the source-free Einstein equations. For other cases the results are not as nice.⁷⁷

The results mentioned above demonstrate that substantive general covariance (in the sense that diffeomorphism invariance is a gauge symmetry) is compatible with having a well-posed initial value problem. But there is clearly a tension between the two, and so one can wonder just how tightly the combination of these two requirements constrains possible laws.⁷⁸

6.4 Cosmic censorship and chronology protection

The positive results reported in the preceding section hardly exhaust the issue of determinism in GTR. One key concern is what happens when the maximal Cauchy development \mathcal{M}, g_{ab} of initial data (Σ, h_{ab}, k_{ab}) satisfying the constraint equations is not maximal simpliciter, i.e. when \mathcal{M}, g_{ab} can be imbedded isometrically as a proper subset of a larger spacetime \mathcal{M}', g'_{ab} satisfying the source-free EFE. The analogous issue can also be raised for the case when $T_{ab} \neq 0$. The future boundary $H^+(\Sigma)$ of the (image of) the future domain of dependence $D^+(\Sigma)$ in the larger spacetime is called the *future Cauchy horizon* of Σ ; the *past Cauchy horizon* $H^-(\Sigma)$ of Σ is defined analogously.⁷⁹ Intuitively, beyond the Cauchy horizons of Σ lie the regions of spacetime where the state of things is not uniquely fixed by the given initial data on Σ ; for generally if the maximal Cauchy development \mathcal{M}, g_{ab} of the initial data is not maximal simpliciter, then the larger extensions for which Σ is not a Cauchy surface are not unique (even up-to-diffeomorphism).

A relatively uninteresting reason why the maximal Cauchy development might be non-maximal simpliciter is that Σ was a poor choice of initial value hypersurface. A trivial but useful example is given by choosing Σ to be the spacelike hyperboloid of Minkowski spacetime pictured in Fig. 3. Here $H^+(\Sigma)$ is the past null cone of the point p .

Some features of this example generalize; in particular, $H^+(\Sigma)$ is always a null surface generated by null geodesics. The unfortunate case can be excluded by requiring, say, that Σ be compact or that it be asymptotically flat. Of course, these conditions exclude many cases of physical relevance, but for sake of discussion let us leave them in place. Even so, the maximal Cauchy development may fail to be maximal simpliciter for more interesting and more disturbing reasons.

⁷⁷For comprehensive reviews of what is known, see [Friedrich and Rendall, 2000] and [Rendall, 2002].

⁷⁸An analysis of gauge symmetries different from the one advertised here is given in [Geroch, 2004]. He gives only two examples of laws that have an initial value formulation and that have diffeomorphism invariance as a gauge symmetry (in his sense).

⁷⁹More precisely, $H^+(\Sigma) := \overline{D^+(\Sigma)} - I^-(D^+(\Sigma))$, and analogously for $H^-(\Sigma)$.




Figure 3. A poor choice of initial value hypersurface

One reason is that a spacetime can start with nice causal properties and evolve in such a way that these properties are lost. The point is illustrated by Misner's $(1 + 1)$ -dim spacetime that captures some of the causal features of Taub-NUT spacetime, which is a solution to the source-free EFE. The Σ in Fig. 4 is a compact spacelike slice in the causally well behaved Taub portion of the spacetime, and its future Cauchy horizon $H^+(\Sigma)$ is a closed null curve. Crossing over this horizon takes one into a region of spacetime where there are closed timelike curves.




Figure 4. Misner spacetime

Another reason that the maximal Cauchy development may fail to be maximal simpliciter is illustrated in Fig. 5 which shows a non-compact asymptotically flat

spacelike slice Σ on which a spherically symmetric ball of matter starts to undergo gravitational collapse. After a finite time the density of collapsing matter becomes infinite, creating a curvature singularity that is pictured as a timelike line in the figure. Strictly speaking, however, it makes no sense to call the singularity a timelike line since the singularity is not part of the spacetime.⁸⁰ But this makes no difference to the main point of relevance here; namely, a causal curve that terminates at a point to the future of $H^+(\Sigma)$ and that is extended into the past may fail to reach Σ , not because it has a past endpoint or because it gets trapped on $H^+(\Sigma)$ (as can happen in the spacetime of Fig. 4) but because it “runs into a singularity” or, better (since the singularity is not part of the spacetime), because it “runs off the edge of the spacetime.”




Figure 5. The development of a naked singularity in spherical gravitational collapse

It is known that EFE plus the imposition of the various energy conditions on T_{ab} discussed above do not suffice to prevent the kind of pathologies indicated by Figs. 4 and 5. But in all of such known examples there is something suspicious: either the matter fields involved are not “fundamental,” i.e. even when gravity is turned off these matter fields are not well behaved in the sense that in Minkowski spacetime the initial value problem for their equations of motion do not admit global existence and uniqueness results (see Section 4.2), or else the initial conditions that eventuate in the singularity are very special, e.g. the initial configuration of matter in Fig. 5 required to produce the singularity must be perfectly spherically symmetric. One might conjecture that what holds for these examples holds generally: Consider a fundamental matter field that can serve as a

⁸⁰One could try to attach the singular points as boundary points of the spacetime manifold. However, the extant prescriptions for doing this lead to counterintuitive features, e.g. the singular points need not be Hausdorff separated from interior points of the manifold; see [Geroch *et al.*, 1982].

source for gravitation. Then the subset of initial data for the Einstein-matter field equations for which the unique (up to diffeomorphism) maximal Cauchy development is not maximal simpliciter is of measure zero in the full space of such data, assuming poor choices of initial value hypersurfaces are excluded. To make this vague claim into a precise conjecture would require a specification of what matter fields are to be counted as fundamental, a specification of a suitable measure on the space of initial data, and a non-question begging specification of what counts as a poor choice of initial value hypersurface. The aimed-for conjecture is referred to as Penrose's *cosmic censorship conjecture*.

Less sweeping versions of the conjecture might focus specifically on one or another of the two types of pathologies illustrated in Figs. 4 and 5. Hawking's *chronology protection conjecture* aims to show that cases where closed timelike curves develop from a causally innocent past are highly non-generic among solutions to Einstein-fundamental-matter-field equations. The *weak cosmic censorship conjecture* aims to show that in generic solutions with asymptotically flat spacetimes, singularities are not "naked" in the sense of being visible to observers at infinity because whatever singularities develop (say, in gravitational collapse) are hidden inside of the event horizons of black holes which serve as one-way causal membranes that shield external observers from any of the pathologies of indeterminism that might develop within the horizon. Some progress has been made in formulating and proving precise versions of chronology protection and weak cosmic censorship, but the juries are still out on strong cosmic censorship.⁸¹

6.5 Predictability in general relativistic spacetimes

In Section 4.3 it was seen that the structure of Minkowski spacetimes has a double-edged quality with respect to determinism and predictability: while this structure makes possible clean examples of determinism, it also makes it impossible for embodied observers who must gather their information about initial conditions by means of causal interactions with the world to use determinism to perform genuine predictions. The point was formalized by defining the domain of predictability $P(q)$ of a point $q \in \mathcal{M}$ of a spacetime \mathcal{M}, g_{ab} and noting that in Minkowski spacetime $P(q) = \emptyset$ for every q . Non-empty domains of predictability are obtained in the modified version of Minkowski spacetime with compactified space slices illustrated in Fig. 2. A feature of this case generalizes to arbitrary general relativistic spacetimes; namely, if the spacetime \mathcal{M}, g_{ab} possesses a Cauchy surface Σ such that $\Sigma \subset I^-(q)$ for some $q \in \mathcal{M}$, then Σ is compact. Since a spacetime with a Cauchy surface Σ is diffeomorphically $\Sigma \times \mathbb{R}$, the kind of complete predictability that comes with having $\Sigma \subset I^-(q)$ for some q is possible only in a spatially finite universe. The converse is not true: the existence of a compact Cauchy surface does not guarantee that there is a Cauchy surface Σ such that $\Sigma \subset I^-(q)$ for some q , de

⁸¹For an overview of progress on cosmic censorship, see [Chruściel, 1992; Isenberg, 1992; [Penrose, 1998; Wald, 1998]. And for an overview of progress on chronology protection, see [Earman *et al.*, 2005].

Sitter spacetime providing a relevant counterexample. Many interesting features of predictability in general relativistic spacetime are studied in [Hogarth, 1993].

6.6 *Determinism and computability in general relativistic spacetimes*

A *Plato machine* for gaining mathematical knowledge about an unresolved conjecture of number theory, whose prenex normal form is $(\forall n_1)(\forall n_2)\dots(\forall n_m)F(n_1, n_2, \dots, n_m)$ or $(\exists n_1)(\exists n_2)\dots(\exists n_m)F(n_1, n_2, \dots, n_m)$ with F recursive, can be conceptualized as an ordinary Turing machine run in Zeno fashion: enumerate the m -tuples of natural numbers and have the computer check in the first 1/2 minute whether F holds of the first tuple, check in the next 1/4 minute whether F holds of the second tuple, etc. At the end of the minute the truth of the conjecture is settled. Despite various claims to the contrary, I see no conceptual incoherency in such a device. But STR militates against the physical instantiation of such a device since the Zeno speed up would seem to require that some of the parts of the device must eventually move faster than the speed of light.⁸²

General relativistic spacetimes seem to open the possibility of creating the functional equivalent of a Plato machine without Zeno tricks and without running afoul of the prohibition on superluminal propagation. Consider a spacetime with the following features. First, there is a timelike half-curve γ_1 with past endpoint, no future endpoint, and an infinite proper length. Second, there is another timelike half-curve γ_2 with past endpoint p and a point $q \in \gamma_2$ such that the proper time elapsed along γ_2 from p to q is finite and such that $\gamma_1 \in I^-(q)$. Such a spacetime has been dubbed a *Malament-Hogarth spacetime*. The theorems of any recursively axiomatizable theory — say, Zermelo-Frankel set theory — can be recursively enumerated, and a device whose worldline is γ_1 can utilize a Turing machine to effectively check each of these theorems to see one has the form “ $0 = 1$ ”. The device can be programmed to send out a signal — “Eureka!” — to an observer whose world line is γ_2 just in case “ $0 = 1$ ” is found among the theorems. Assuming that the observer γ_2 is aware of this arrangement, she gains knowledge of the consistency/inconsistency of ZF: she knows that ZF is consistent just in case she has not received a “Eureka!” signal by the time she reaches the point q .

Similar arrangements can be used to “decide,” at least in principle, Turing undecidable questions and to “compute” Turing uncomputable functions (see [Hogarth, 1994]). They, thus, threaten to falsify the physical Church-Turing thesis which asserts, roughly, that any physical computing device can be simulated by a Turing machine (see [Etsefi and Németi, 2002] for a careful formulation of this thesis). In contrast to the original Church-Turing thesis which belongs to mathematical logic, the physical Church-Turing thesis lies in the borderland of mathematical logic and physics (see [Deutsch *et al.*, 2000]), and it is much harder to evaluate, especially if it is understood to require the physical realizability of the devices that implement the bifurcated supertask. Here I will confine myself to a few remarks

⁸²Perhaps conflict with STR can be avoided by Zeno shrinking the parts, but this maneuver may run afoul of quantum restrictions.

on this matter and refer the interested reader to Németi and David [2005] for a fuller discussion.

Malament-Hogarth spacetimes are among the solutions of EFE — e.g. Reissner-Nordström spacetime and (the universal covering spacetime of) anti-De Sitter spacetimes. These particular spacetimes do not involve causal anomalies in the sense that they admit global time functions. However, all Malament-Hogarth spacetimes fail to be globally hyperbolic. Indeed, it can be shown of such spacetimes that if $\Sigma \subset \mathcal{M}$ is any spacelike hypersurface such that the above defined γ_1 lies in $I^+(\Sigma)$, then any Malament-Hogarth point q whose chronological past contains γ_1 must lie on or beyond $H^+(\Sigma)$ (see Lemma 4.3 of [Earman, 1995, 117]). The possibility of non-deterministic influences, which might open the possibility that γ_1 receives a false “Eureka!” message, seems to undermine the use of Malament-Hogarth spacetimes for gaining knowledge in the sense of certainty. However, one should not draw hasty conclusions here since, as discussed in the following subsection, it is possible to have deterministic dynamics for fields propagating on a non-globally hyperbolic spacetime. Also it might seem that the problem is avoided by the fact that it can be arranged so that any signal from γ_1 arrives at γ_2 before the Malament-Hogarth point q and, thus, within $D^+(\Sigma)$. But since a “Eureka!” message can arrive arbitrarily close to q , the receiver must possess arbitrarily accurate discriminatory powers to separate signals that arrive before q from the potentially false signals that arrive after q .

6.7 *The possibility of deterministic dynamics in non-globally hyperbolic spacetimes*

For fields that propagate on a general relativistic spacetime, the failure of global hyperbolicity can undermine the initial value problem. For example, it is known that in generic two-dimensional spacetimes with closed timelike curves (CTCs) the scalar wave equation may fail to have smooth solutions or else may admit multiple solutions for the same initial data specified on a spacelike hypersurface. But remarkably, existence and uniqueness results have been proven for some four-dimensional spacetimes with CTCs (see [Friedman, 2004] for a review).

For spacetimes that do not have such blatant causal anomalies as CTCs but which nevertheless fail to be globally hyperbolic, Hilbert space techniques can sometimes be used to cure breakdowns in existence and uniqueness.⁸³ Consider a general relativistic spacetime \mathcal{M}, g_{ab} that is static and possesses a global time function. The first condition means that there is a timelike Killing field V^a that is hypersurface orthogonal.⁸⁴ The second condition can be guaranteed by choosing

⁸³The use of Hilbert space techniques to study problems in classical physics was pioneered by Koopman [1931]. However, Koopman’s approach assumes determinism at the classical level and then shows how to represent this determinism as unitary evolution on a Hilbert space.

⁸⁴The Killing condition is $\nabla_{(c}g_{ab)} = 0$ where ∇_a is the covariant derivative operator compatible with g_{ab} . Staticity guarantees that locally a local coordinate system (x^α, t) can be chosen so that the line element takes the form $ds^2 = g_{\alpha\beta}(x^\gamma)dx^\alpha dx^\beta - g_{44}(x^\gamma)dt^2$. Cf. Malament, this volume, section 2.7.

a spacelike hypersurface Σ orthogonal to V^a and by requiring that every integral curve of V^a meets Σ in exactly one point. Then every point $p \in \mathcal{M}$ can be labeled by (x, t) , where $x \in \Sigma$ is the point where the integral curve of V^a through p meets Σ , and t is the value of the Killing parameter along this integral curve. Such a causally well behaved spacetime can nevertheless fail to be globally hyperbolic because, intuitively speaking, it possess a naked, timelike singularity. (To help fix intuitions, think of Minkowski spacetime with a timelike world tube removed. Or the reader familiar with GTR can think of the negative mass Schwarzschild solution to EFE, which is static and contains a timelike naked singularity at $r = 0$.) Now consider a massive $m \geq 0$ scalar field ϕ propagating on this background spacetime in accord with the Klein-Gordon equation (13). For the type of spacetime in question this equation can be rewritten in the form

$$(17) \quad \frac{\partial^2 \phi}{\partial t^2} = -A\phi$$

where t is the Killing parameter (see [Wald, 1980a], [Horowitz and Marolf, 1995]). The differential operator A can be considered a Hilbert space \hat{A} operator acting on $L^2_{\mathbb{R}}(\Sigma, d\vartheta)$, where $d\vartheta$ is the volume element of Σ divided by $\sqrt{-V^a V_a}$. With the domain initially taken to be $C_0^\infty(\Sigma)$, \hat{A} is a positive symmetric operator. The proposal is to replace the partial differential equation (18) with the ordinary differential equation

$$(18) \quad \frac{d^2 \phi}{dt^2} = -\hat{A}\phi$$

where the time derivative in (19) is a Hilbert space derivative. Since the Hilbert space operator \hat{A} is real it has self-adjoint extensions, and since \hat{A} is positive the positive square root of a self-adjoint extension \hat{A}_e can be extracted to give

$$(19) \quad \phi(t) := \cos(\sqrt{\hat{A}_e}t)\phi(0) + \sin(\sqrt{\hat{A}_e}t)\dot{\phi}(0)$$

which is valid for all t and all $\phi(0), \dot{\phi}(0) \in \mathcal{H}$. Since $\phi(t)$ is a solution throughout the spacetime of the Klein-Gordon equation (13) and since it is the unique solution for given initial data $\phi(0), \dot{\phi}(0)$ on Σ , it provides (relative to the chosen self-adjoint extension) a deterministic prescription for the dynamics of the Klein-Gordon field. There are other possible prescriptions for obtaining the dynamics of ϕ , but Ishibashi and Wald [2003] have shown that the one just reviewed is the only one satisfying the following set of restrictions: it agrees locally with (18); it admits a suitable conserved energy; it propagates the field causally; and it obeys time translation and time reflection invariance. If the Hilbert space operator \hat{A} is essentially self-adjoint, then the unique self-adjoint extension \hat{A}_e provides *the* dynamics for the ϕ field satisfying the said restrictions. And this dynamics is fully deterministic despite the fact that the background spacetime on which the field propagates is not globally hyperbolic. Not surprisingly, however, \hat{A} fails to be essentially self-adjoint for many examples of static but non-globally hyperbolic spacetimes, and

unless further restrictions are added to single out one of the self-adjoint extensions, no unambiguous dynamics is specified by the above procedure. But remarkably, Horowitz and Marolf [1995] have provided examples of static, non-globally hyperbolic spacetimes where \hat{A} is essentially self-adjoint, and in these cases the above prescription produces a dynamics of the ϕ field that is fully deterministic despite the presence of naked singularities.

7 DETERMINISM IN RELATIVISTIC QFT

Ordinary QM starts from a classical mechanical description of a system of particles — specifically, a Hamiltonian description — and attempts to produce a quantized version. Similarly, QFT starts from a classical relativistic description of a field and attempts to produce a quantized version. However, some classical fields do not lend themselves to a QFT that physicists find acceptable. Consider, for example, the non-linear wave equation (13) as a candidate for describing boson-boson interactions. A heuristic quantization procedure leads to the conclusion that there is no lowest energy state, leaving the system vulnerable to radiative collapse. On these grounds quantum field theorists have categorized the hypothetical interaction as “not physically realizable” (see [Baym, 1960]). That difficulties are encountered in QFT is perhaps not surprising when it is realized that the field in question is ill-behaved at the classical level in that regular initial data can pick out solutions that develop singularities within a finite amount of time. Is it plausible that deterministic behavior at the classical relativistic level can serve as a selection principle for what fields it is appropriate to quantize?

Determinism also plays a more constructive role in QFT. In ordinary QM, quantization involves the choice of a suitable representation of the canonical commutation relations $[\hat{x}_j, \hat{p}_k] = i\delta_{jk}$ (CCR). Since unbounded operators are involved, this form of the CCR only makes sense when the domains of the operators are specified. Such worries can be avoided by working with the Weyl, or exponentiated, form of the CCR, which also makes available the Stone-von Neumann theorem: for a finite number of degrees of freedom, the irreducible strongly continuous representations of the Weyl CCR are all unitarily equivalent — in fact, all are equivalent to the familiar Schrödinger representation. This theorem no longer applies when there are an infinite number of degrees of freedom, as in QFT, a feature of QFT that raises a number of interesting interpretational issues that are not relevant here. What is relevant is the fact that the construction of the CCR algebra for, say, the Klein-Gordon field in Minkowski spacetime, makes essential use of the deterministic propagation of this field (see [Wald, 1994]). This construction can be generalized to a Klein-Gordon field propagating in an arbitrary general relativistic background spacetime that is globally hyperbolic since the deterministic nature of the propagation carries over to the more general setting.

For a non-globally hyperbolic spacetime \mathcal{M}, g_{ab} it is still the case that for any $p \in \mathcal{M}$ there is a neighborhood $\mathcal{N}(p)$ such that $\mathcal{N}, g_{ab}|_{\mathcal{N}}$, considered as a spacetime in its own right, is globally hyperbolic, and thus the field algebra $\mathcal{A}(\mathcal{N})$ for this

mini-spacetime can be constructed by the usual means. One can then ask whether these local algebras can be fitted together to form a global algebra $\mathcal{A}(\mathcal{M})$ with the natural net properties (e.g. each such $\mathcal{A}(\mathcal{N})$ is a subalgebra of $\mathcal{A}(\mathcal{M})$, and if $\mathcal{N}_1 \subset \mathcal{N}_2$ then $\mathcal{A}(\mathcal{N}_1)$ is a subalgebra of $\mathcal{A}(\mathcal{N}_2)$). Kay (1992) calls the spacetimes for which the answer is affirmative *quantum compatible*, the idea being that non-quantum compatible spacetimes are not suitable arenas for QFT. A variety of non-globally hyperbolic spacetimes are not quantum compatible, e.g. 2-dim cylindrical spacetimes obtained from two-dimensional Minkowski spacetime by identifications along the time axis. But, remarkably, some acausal spacetimes have been shown to be quantum compatible (see [Fewster and Higuchi, 1996] and [Fewster, 1999]).

8 DETERMINISM AND QUANTUM GRAVITY

Arguably the biggest challenge in theoretical physics today is to combine the insights of GTR and QFT so as to produce a quantum theory of gravity (see [Rovelli, this vol.]). Some inkling of what this sought after theory will yield can perhaps be gained from semi-classical quantum gravity, which is a kind of shot-gun marriage of GTR and QFT. Semi-classical means that there is no attempt to quantize the metric of spacetime, but rather than merely treating a general relativistic spacetime as a fixed background on which quantum fields propagate (as in the preceding section), an attempt is made to calculate the back-reaction on the metric by inserting the quantum expectation value of the (renormalized) stress-energy in place of the classical stress-energy tensor on the right hand side of EFE (17). Although there may be no consistent theory underlying such a procedure, good theoretical physicists know how to extract usable information from it. Perhaps the most spectacular extraction is Hawking's conclusion that a black hole is not black but radiates exactly like a black body at temperature proportional to the surface gravity of the black hole. This *Hawking effect* is taken as confirmation that the formula for black hole entropy,⁸⁵ which had been derived by independent means, is more than a formal expression; it shows that black hole entropy is the ordinary thermodynamic entropy of a black hole (see [Wald, 1994]).⁸⁶ Theoretical physicists of different schools are in agreement that this is a stable result that has to be accommodated by an adequate quantum theory of gravity. But from this point on, the disagreements increase to the point of stridency.

The Hawking effect means that, when quantum effects are taken into account, black holes are not stable objects because the Hawking radiation must be accompanied by a diminution of the mass of the black hole. Presumably, as this process

⁸⁵ $S_{bh} = \frac{kc^3}{4G\hbar} A$, where A is the surface area of the black hole.

⁸⁶ The Hawking effect is related to, but distinct from, the *Unruh effect*. The latter effect is analyzed in terms of the apparatus of quantum statistical mechanics discussed in [Emch, this vol.]. In Minkowski spacetime the essence of the Unruh effect is that what an observer uniformly accelerated through the Minkowski vacuum experiences is described by a KMS state. The Unruh effect has been generalized to general relativistic spacetime; see [Kay and Wald, 1991].

goes deeper and deeper into the quantum regime, the semi-classical calculation will eventually break down. But *if* the continuation of the calculation can be trusted, then in the fullness of time the black hole will completely evaporate. (The estimated evaporation time for a black hole of solar mass is the order of 10^{67} years, much greater than the age of the universe. But this is no problem in a universe with an infinite future, as the latest cosmological measurements indicate is the case for our universe.) And *if* the result of the evaporation can be described by a classical general relativistic spacetime, the result is a momentarily naked singularity and a breakdown in global hyperbolicity, as is indicated in Fig. 6.⁸⁷ So even if some form of cosmic censorship holds for classical GTR, quantum effects seem to undo it.




Figure 6. Conformal representation of black hole evaporation

Having gotten this far it is not difficult to establish that if at a time Σ_1 prior to the evaporation of the black hole the quantum field is in a pure state and if Hawking radiation has established correlations between relatively spacelike regions, such as the region R_1 in the black hole interior (see Fig. 6) and the region consisting of a “sandwich” about the post evaporation time Σ_2 , then the state of the quantum field will be mixed at a post evaporation time Σ_2 .⁸⁸ Since a pure-to-mixed state transition is necessarily non-unitary, the upshot is a loss of unitarity.⁸⁹

⁸⁷Following the conventions of conformal diagrams (see [Hawking and Ellis, 1973]), \mathfrak{I}^+ denotes future null infinity (the terminus of outgoing null rays), and ι^0 denotes spatial infinity.

⁸⁸This can be rigorously established in the algebraic formulation of QFT; see [Earman, 2002].

⁸⁹And, incidentally, there is also a violation of time reversal invariance; see [Wald, 1986] and [Earman, 2002].

This “information loss paradox,” as it is often referred to in the physics and the popular literature, has evoked an amazing variety of reactions; see [Belot *et al.*, 1999] for an overview. Most notable are the reactions from those who are so desperate to avoid the conclusion that they are willing to deploy “black hole complementarity”⁹⁰ and thereby abandon the mainstream reading of relativity theory, namely, that what the theory teaches us is that there is an intrinsic observer-independent reality — the very opposite of vulgar relativism that has it that everything is relative-to-an-observer.

But stepping back from the fray allows one to see that there is no need for such desperate measures. The pure-to-mixed evolution that is at the heart of the “paradox” need not be seen as a breakdown of quantum theory.⁹¹ Nor is it surprising that consequences labeled ‘paradoxical’ flow from loss of global hyperbolicity. What needs to be questioned is whether this loss of global hyperbolicity is a plausible expectation of quantum gravity. Semi-classical quantum gravity suggests such a loss, but this way of bringing GTR and QFT together is at best a stepping stone to a full theory of quantum gravity. And just as ordinary QM showed the ability to smooth away singularities of classical mechanics, so the correct theory of quantum gravity may show the ability to smooth away the singularities of classical GTR.

Some positive indications come from the work of string theorists who are able to point to mechanisms that can smooth out singularities in classical general relativistic models; for example, Johnson *et al.* [2000] show that brane repulsion smooths out a class of naked singularities dubbed the *repulsion*. String theorists can also give a back door argument for exclusion of some types of classical singularities: postulate or prove that the sought after M-theory gives a stable ground state, and then note that this rules out analogues of the negative mass Schwarzschild solution and the like.

Other encouraging results come from loop quantum gravity (LQG), which aims to produce a quantum theory of gravity by applying to GTR a version of the canonical quantization based on a new set of canonical variables introduced by Amitaba Sen and exploited by Abay Ashtekar.⁹² In the Friedmann-Robertson-Walker big bang models of classical GTR the scale factor a of spacetime goes to zero as the big bang singularity is approached, and the curvature blows up since

⁹⁰Consider the case in STR of two inertial observers, O and O' , who describe an ambient electromagnetic field using electric and magnetic fields (\mathbf{E}, \mathbf{B}) and $(\mathbf{E}', \mathbf{B}')$ respectively. There is a translation between the two descriptions which gives \mathbf{E}' and \mathbf{B}' as functions of \mathbf{E} , \mathbf{B} , and the relative velocity of O and O' and vice versa with O and O' exchanged. The existence of such a translation follows from the fact that there is an intrinsic, observer independent reality — in this case, the electromagnetic field as specified by the Maxwell tensor field. This tensor field is independent of coordinate systems, reference, frame, and observers. Contracting it with different velocity fields, representing the motions of different observers, results in different descriptions in terms of electric and magnetic fields. Whatever else it means, the “complementarity” part of “black hole complementarity” means that the different descriptions of an evaporating black hole given by two observers, one who falls through the black hole horizon and one who remains outside the horizon, are not related in the way the descriptions of O and O' are related.

⁹¹See [Wald, 1994, 181–182] and [Belot *et al.*, 1999].

⁹²See [Rovelli, 2004] and [this vol.] for surveys of loop quantum gravity.

it scales as $1/a^2$.⁹³ Since there is no physically motivated way to extend such a solution through the initial singularity, the question of what happens “before” the big bang belongs to theology or science fiction rather than science. The situation is startlingly different in LQG. Corresponding to the classical quantity $1/a$ there is a self-adjoint operator, acting on the Hilbert space of spatially homogeneous, isotropic quantum kinematical states, and its spectrum is bounded from above, giving a first indication that the classical singularity has been removed (see [Bojowald, 2001]). A complete proof of removal would require that the quantum dynamics gives an unambiguous evolution through the classical singularity. In LQG the “dynamics” is obtained by solving the Hamiltonian constraint equation, which restricts the physically allowed states. For the case at issue this constraint equation comes in the form of a difference equation rather than a differential equation. If the scale factor a is regarded as a “clock variable,” then the constraint equation provides a “time evolution” of the quantum state through discrete steps of the clock variable. The key point is that this evolution equation does determine a unique continuation through the classical singularity.⁹⁴ However, what happens at the classical singularity is undetermined because the coefficient corresponding to this stage decouples from the other coefficients in the evolution equation (see [Ashtekar and Bojowald, 2003] for details).

From the point of view of determinism this last result means that the situation is somewhat ironic. Determinism is not threatened in classical GTR by the initial big bang singularity of the Friedmann-Robertson-Walker models because these models are globally hyperbolic, and because there is no physically motivated way to extend through the initial singularity. In LQG the initial singularity is banished both in the sense that curvature remains bounded and in the sense that there is a sensible way to extend through the classical singularity. But the price to be paid is a loss of determinism in LQG at the classical singularity, which can be seen as a Cheshire grin of the classical singularity.

Recently LQG has been used to resolve black hole singularities, leading to a new perspective on the Hawking information loss paradox in which Fig. 6 is not a valid depiction of black hole evaporation (see [Ashtekar *et al.*, 2005]). It is argued that, analogously to the FRW case, the quantum evolution continues through the classical singularity.⁹⁵ The new picture is not one in which global hyperbolicity is restored; indeed, that concept is not meaningful since what replaces the classical singularity is a region which cannot be described even approximately by the spacetime geometry of classical GTR. Nevertheless, it is argued that in the quantum

⁹³The line element of FRW models can be written in the form $ds^2 = a(t)d\sigma^2 - dt^2$, where $d\sigma^2$ is the spatial line element.

⁹⁴But see [Green and Unruh, 2004] where it is shown that in a spatially closed FRW model, the use of the scale factor as a “clock variable” is problematic. And the situation in inhomogeneous cosmologies is much more delicate and complicated; see [Brunnemann and Thiemann, 2006a; 2006b].

⁹⁵As in the FRW case, the Hamiltonian constraint equation becomes a difference equation. The “quantum evolution” comes from this equation by choosing a suitable “clock variable” and then following the quantum state through discrete steps of the clock variable.

evolution a pure state remains pure and, in this sense, no information is lost. In its present form the argument has a heuristic character, and detailed calculations are needed to make it rigorous.

9 CONCLUSION

Is the world deterministic? Without the aid of metaphysical revelation, the only way we have to tackle this question is to examine the fruits of scientific theorizing. We can thus set ourselves the task of going through the theories of modern physics and asking for each: If the world is the way it would have to be in order for the theory to be true, is it deterministic? One of the things we discovered is that this task is far from straightforward, for the way in which theories are interpreted is colored by our attitudes towards determinism. For example, the unwillingness to see determinism fail at the starting gate in Newtonian gravitational theory militates in favor of taking gravitation to be a direct interparticle interaction and against assigning independent degrees of freedom to the Newtonian gravitational field. And an unwillingness to see determinism fail at the starting gate in GTR leads to the rejection of a naively realistic interpretation of the textbook version of the theory's description of spacetime and to the acceptance of diffeomorphism invariance as a gauge symmetry — which entails that none of the variables used in textbook presentations is counted as a genuine (= gauge invariant) physical magnitude.

The fortunes of determinism are too complicated to admit of a summary that is both short and accurate, but roughly speaking the story for classical (= non-quantum theories) is this. In Newtonian theories determinism is hard to achieve without the aid of various supplementary assumptions that threaten to become question-begging. For special relativistic theories determinism appears so secure that it is used as a selection criterion for “fundamental fields.” GTR, under the appropriate gauge interpretation, is deterministic locally in time; but whether it is deterministic non-locally in time devolves into the unsettled issues of cosmic censorship and chronology protection.

Quantum physics is the strangest and most difficult case. Ordinary QM is in some respects more deterministic than Newtonian mechanics; for example, QM is able to cure some of the failures of Newtonian determinism which occur either because of non-uniqueness of solutions or the breakdown of solutions. But the fortunes of determinism in QM ultimately ride on unresolved interpretational issues. The main driving force behind these issues is the need to explain how QM can account for definite outcomes of experiments or more generally, the apparent definiteness of the classical world — an ironic situation since QM is the most accurate physical theory yet devised. Some of the extant responses to this explanatory challenge would bury determinism while others give it new life.

A new arena for testing the mettle of determinism is provided by the nascent quantum theories of gravity. There are some preliminary indications that just as ordinary QM was able to smooth out singularities of Newtonian mechanics,

so quantum gravity effects may smooth out singularities of classical GTR. If this smoothing ability is broad enough it would alleviate worries that there are analogues in quantum gravity of breakdowns in determinism in classical GTR associated with failures of cosmic censorship. Quantum gravity will likely put a new face on the measurement problem and related interpretational issues that arise in ordinary QM. But it is too early to say whether this new face will smile on determinism.

ACKNOWLEDGMENT

It is a pleasure to acknowledge the help and advice I have received from Gordon Belot, Michael Dickson, Klaas Landsman, Carlo Rovelli, and Laura Ruetsche. Special thanks are due to Jeremy Butterfield for numerous improvements. But, needless to say, there should be no presumption that any of the above mentioned share the opinions expressed herein.

BIBLIOGRAPHY

- [Albert, 1992] D. Albert. *Quantum Mechanics and Experience*. Cambridge, MA: Harvard University Press, 1992.
- [Albert, 2000] D. Albert. *Time and Chance*. Cambridge, MA: Harvard University Press, 2000.
- [Anderson and Thompson, 1992] I. Anderson and G. Thompson. *The inverse problem of the calculus of variations for ordinary differential equations. Memoirs of the American Mathematical Society*, No. 473. Providence, RI: American Mathematical Society, 1992.
- [Ashtekar and Bojowald, 2005] A. Ashtekar and M. Bojowald. Black hole evaporation: a paradigm. 2005. *gr-qc/0504029*.
- [Ashtekar et al., 2003] A. Ashtekar, M. Bojowald, and J. Lewandowski. Mathematical structure of loop quantum cosmology. *Advances in Theoretical and Mathematical Physics*, 7: 233–268, 2003. *gr-qc/0304074*.
- [Atmanspacher and Bishop, 2002] H. Atmanspacher and R. Bishop (eds.). *Between Chance and Choice: Interdisciplinary Perspectives on Determinism*. Thorveton: Imprint Academic, 2002.
- [Ax and Kochen, 1999] J. Ax and S. Kochen. Extension of quantum mechanics to individual systems. 1999. *quant-ph/9905077*.
- [Bacciagaluppi and Dickson, 1998] G. Bacciagaluppi and M. Dickson. Dynamics for modal interpretations. *Foundations of Physics*, 29: 1165–1201, 1998.
- [Barbour, 1974] J. B. Barbour. Relative-distance Machian theories. *Nature*, 249: 328–329, 1974.
- [Barbour, 1999] J. B. Barbour. The development of Machian themes in the twentieth century. In J. Butterfield (ed.), *The Arguments of Time*, pages 83–110. Oxford: Oxford University Press, 1999.
- [Barbour and Bertotti, 1977] J. B. Barbour and B. Bertotti. Gravity and inertia in a Machian framework. *Nuovo Cimento*, 38B: 1–27, 1977.
- [Bargmann, 1954] V. Bargmann. On unitary ray representations of continuous groups. *Annals of Mathematics*, 59: 1–46, 1954.
- [Barrett, 1999] J. A. Barrett. *The Quantum Mechanics of Minds and Worlds*. Oxford: Oxford University Press, 1999.
- [Baym, 1960] G. Baym. Inconsistency of cubic boson-boson interactions. *Physical Review*, 117: 886–888, 1960.
- [Beig, 2005] R. Beig. Concepts of hyperbolicity and relativistic continuum mechanics. 2005. *gr-qc/0411092*.
- [Belot and Earman, 1997] G. Belot and J. Earman. Chaos out of order: quantum mechanics, the correspondence principle, and chaos. *Studies in History and Philosophy of Modern Physics*, 28B: 147–182, 1997.

- [Belot and Earman, 1999] G. Belot and J. Earman. From metaphysics to physics. In J. Butterfield and C. Pagonis (eds.), *From Physics to Philosophy*, pages 166–186. Cambridge: Cambridge University Press, 1999.
- [Belot *et al.*, 1999] G. Belot, J. Earman, and L. Ruetsche. The Hawking information loss paradox: The anatomy of a controversy. *British Journal for the Philosophy of Science*, 50: 189–229, 1999.
- [Bojowald, 2001] M. Bojowald. Absence of a singularity in loop quantum cosmology. *Physical Review Letters*, 86: 5227–5230, 2001.
- [Brading and Brown, 2003] C. Brading and H. Brown. Symmetries and Noether’s theorems. In C. Brading and H. Brown (eds.), *Symmetries in Physics: Philosophical Reflections*, pages 89–109. Cambridge: Cambridge University Press, 2003.
- [Brown, 1999] H. R. Brown. Aspects of objectivity in quantum mechanics. In J. Butterfield and C. Pagonis (eds.), *From Physics to Philosophy*, pages 45–70. Cambridge: Cambridge University Press, 1999.
- [Brunneman and Thieman, 2006a] J. Brunneman and T. Thiemann. On (cosmological) singularity avoidance in loop quantum gravity, *Classical and Quantum Gravity*, 23: 1395–1427, 2006.
- [Brunneman and Thieman, 2006b] J. Brunneman and T. Thiemann. Unboundedness of triad-like operators in loop quantum gravity, *Classical and Quantum Gravity*, 23: 1429–1483, 2006.
- [Bub, 1997] J. Bub. *Interpreting the Quantum World*. Cambridge: Cambridge University Press, 1997.
- [Butterfield, 1998] J. Butterfield. Determinism and indeterminism. *Routledge Encyclopedia of Philosophy*, 3: 33–39. New York: Routledge, 1998.
- [Callender, 2005] C. Callender. Time is the simplest (and strongest) thing. Pre-print, 2005.
- [Carroll, 2004] J. Carroll (ed.). *Readings on Laws of Nature*. Pittsburgh, PA: University of Pittsburgh Press, 2004.
- [Cassirer, 1956] E. Cassirer. *Determinism and Indeterminism in Modern Physics*. New Haven: Yale University Press, 1956.
- [Chalmers, 1970] A. F. Chalmers. Curie’s principle. *British Journal for the Philosophy of Science*, 21: 133–148, 1970.
- [Chruściel, 1992] P. Chruściel. On uniqueness in the large of solutions of Einstein’s equations (‘strong cosmic censorship’). In M. Gotay, J. E. Marsden, and V. Moncrief (eds.), *Mathematical Aspects of Classical Field Theory*, pages 235–273. Providence, RI: American Mathematical Society, 1992.
- [Courant and Hilbert, 1962] R. Courant and D. Hilbert. *Methods of Mathematical Physics*, 2 vols. New York: Interscience Publishers, 1962.
- [Curie, 1894] P. Curie. Sur la symétrie des phénomènes physiques: symétrie d’un champ électrique et d’un champ magnétique. *Journal de Physique*, 3: 393–415, 1894.
- [Cushing, 1994] J. Cushing. *Quantum Mechanics: Historical Contingency and the Copenhagen Hegemony*. Chicago: University of Chicago Press, 1994.
- [Cushing and Bowman, 1999] J. Cushing and G. Bowman. Bohmian mechanics and chaos. In J. Butterfield and C. Pagonis (eds.), *From Physics to Philosophy*, pages 90–107. Cambridge: Cambridge University Press, 1999.
- [Deutsch *et al.*, 2000] D. Deutsch, A. Ekert, and R. Lupacchini. Machines, logic, and quantum physics. *Bulletin of Symbolic Logic*, 6: 265–283, 2000.
- [Dickson, 1997] W. M. Dickson. On the plurality of dynamics: transition probabilities and modal interpretations. In R. A. Healey and G. Hellman (eds.), *Quantum Measurement: Beyond Paradox*. *Minnesota Studies in the Philosophy of Science*, Vol. 17, pages 160–182. Minneapolis, MN: University of Minnesota Press, 1997.
- [Dickson, 1998] W. M. Dickson. *Quantum Chance and Non-Locality: Probability and Non-Locality in the Interpretations of Quantum Mechanics*. Cambridge: Cambridge University Press, 1998.
- [Driver, 1977] R. D. Driver. *Ordinary and Delay Differential Equations*. New York: Springer-Verlag, 1977.
- [Driver, 1979] R. D. Driver. Can the future influence the present? *Physical Review D*, 19: 1098–1107, 1979.
- [Earman, 1989] J. Earman. *World Enough and Space-Time: Absolute vs. Relational Theories of Space and Time*. Cambridge, MA: MIT Press/Bradford Books, 1989.

- [Earman, 1995] J. Earman. *Bangs, Crunches, Whimpers, and Shrieks: Singularities and Acausalities in Relativistic Spacetimes*. Oxford: Oxford University Press, 1995.
- [Earman, 2002] J. Earman. What time reversal invariance is and why it matters. *International Journal for the Philosophy of Science*, 16: 245–264, 2002.
- [Earman, 2003] J. Earman. Tracking down gauge: an ode to the constrained Hamiltonian formalism. In K. Brading and E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*, pages 140–162. Cambridge: Cambridge university Press, 2003.
- [Earman, 2004a] J. Earman. Determinism: what we have learned and what we still don't know. In J. K. Campbell, M. O'Rourke, and D. Shier (eds.), *Determinism, Freedom, and Agency*, pages 21–46. Cambridge, MA: MIT Press, 2004a.
- [Earman, 2004b] J. Earman. Curie's principle and spontaneous symmetry breaking. *International Studies in Philosophy of Science*, 18: 173–199, 2004b.
- [Earman, 2005] J. Earman. Essential self-adjointness: Implications for determinism and the classical-quantum correspondence. Pre-print, 2005b.
- [Earman, 2006] J. Earman. Two challenges to the substantive requirement of general covariance. *Synthese*, 148: 443–468, 2006.
- [Earman and Roberts, 2006] J. Earman and J. Roberts. Contact with the nomic: a challenge to the deniers of humean supervenience. *Philosophy and Phenomenological Research*, in press, 2006.
- [Earman et al., 2005] J. Earman, C. Smeenk, and C. Wüthrich. Do the laws of physics forbid the operation of a time machine? Pre-print, 2005.
- [Earman and Ruetsche, 2005] J. Earman and L. Ruetsche. Relativistic invariance and modal interpretations. *Philosophy of Science*, forthcoming, 2005.
- [Etesi and Némethi, 2002] G. Etesi and I. Némethi. Turing computability and Malament-Hogarth spacetimes. *International Journal of Theoretical Physics*, 41: 342–370, 2002.
- [Fefferman, 2001] C. Fefferman. Existence and smoothness of the Navier-Stokes equation. Clay Mathematics Institute, Millennium Problems, 2001. <http://www.claymath.org/>
- [Fewster, 1999] C. J. Fewster. Bisolutions to the Klein-Gordon equation and quantum field theory on 2-dimensional cylinder spacetimes. *Classical and Quantum Gravity*, 16: 769–788, 1999.
- [Fewster and Higuchi, 1996] C. J. Fewster and A. Higuchi. Quantum field theory on certain non-globally hyperbolic spacetimes. *Classical and Quantum Gravity*, 13: 51–62, 1996.
- [Fine, 1982a] A. Fine. Joint distributions, quantum correlations, and commuting observables. *Journal of Mathematical Physics*, 23: 1306–1310, 1982.
- [Fine, 1982b] A. Fine. Hidden variables, joint probability, and the Bell inequalities. *Physical Review Letters*, 48: 291–295, 1982.
- [Foias et al., 2001] C. Foias, O. Manley, R. Rosa, and R. Teman. *Navier-Stokes Equations and Turbulence*. Cambridge: Cambridge University Press, 2001.
- [Friedrich and Rendall, 2000] H. Friedrich and A. Rendall. The Cauchy problem for the Einstein equations. In B. G. Schmidt (ed.), *Einstein's Field Equations and their Physical Implications: Selected Essays in Honor of Jürgen Ehlers*, pages 127–223. New York: Springer-Verlag, 2000.
- [Friedman, 2004] J. L. Friedman. The Cauchy problem on spacetimes that are not globally hyperbolic. In P. T. Chruściel and H. Friedrich (eds.), *The Einstein Equations and the Large Scale Behavior of Gravitational Fields*, pages 331–346. Basel: Birkhäuser, 2004.
- [Geroch, 1977] R. P. Geroch. Prediction in general relativity. In J. Earman, C. Glymour, and J. Stachel (eds.), *Foundations of Spacetime Theories. Minnesota Studies in the Philosophy of Science*, Vol. 8, pages 81–93. Minneapolis, MN: University of Minnesota Press, 1977.
- [Geroch, 1995] R. P. Geroch. Relativistic theories of dissipative fluids. *Journal of Mathematical Physics*, 36: 4226–4241, 1995.
- [Geroch, 2004] R. P. Geroch. Gauge, diffeomorphisms, initial-value formulation, etc. In P. T. Chruściel and H. Friedrich (eds.), *The Einstein Equations and the Large Scale Behavior of Gravitational Fields*, pages 441–477. Basel: Birkhäuser, 2004.
- [Geroch et al., 1982] R. P. Geroch, C. B. Liang, and R. M. Wald. Singular boundaries of spacetimes. *Journal of Mathematical Physics*, 23: 432–435, 1982.
- [Ghirardi et al., 1986] G. C. Ghirardi, A. Rimini, and T. Weber. Unified dynamics for microscopic and macroscopic systems. *Physical Review D*, 34: 470–491, 1986.
- [Green and Unruh, 2004] D. Green and W. H. Unruh. Difficulties with recollapsing models in closed isotropic loop quantum cosmology. *Physical Review D*, 103502-1-7, 2004.

- [Hadamard, 1923] J. Hadamard. *Lectures on Cauchy's Problem in Linear Partial Differential Equations*, 1923. Reprinted by Dover, New York, 1952.
- [Hawking and Ellis, 1973] S. W. Hawking and G. F. R. Ellis. *The Large Scale Structure of Space-Time*. Cambridge: Cambridge University Press, 1973.
- [Henneaux and Teitelboim, 1992] M. Henneaux and C. Teitelboim. *Quantization of Gauge Systems*. Princeton, NJ: Princeton University Press, 1992.
- [Hogarth, 1993] M. Hogarth. Predicting the future in general relativistic spacetimes. *Studies in History and Philosophy of Science*, 24B: 721–739, 1993.
- [Hofer, 2003] C. Hofer. Causal Determinism. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/sum2005/entries/determinism-causal/>
- [Hogarth, 1994] M. Hogarth. Non-Turing computers and non-Turing computability. In D. Hull, M. Forbes, and R. M. Burian (eds.), *PSA 1994*, Vol. 1, pages 126–138, 1994.
- [Holland, 1993] P. R. Holland. *The Quantum Theory of Motion*. Cambridge: Cambridge University Press, 1993.
- [Horowitz and Marolf, 1995] G. Horowitz and D. Marolf. Quantum probes of spacetimes singularities. *Physical Review D*, 52: 5670–5675, 1995.
- [Howard, 1999] D. Howard. Point coincidences and pointer coincidences: Einstein and the invariant content of space-time theories. In H. Goenner, J. Renn, J. Ritter, and T. Sauer (eds.), *The Expanding Worlds of General Relativity, Einstein Studies*, Vol. 7, pages 463–500. Boston: Birkhäuser, 1999.
- [Isenberg, 1992] J. Isenberg. Progress on strong cosmic censorship. In M. Gotay, J. E. Marsden, and V. Moncrief (eds.), *Mathematical Aspects of Classical Field Theory*, pages 403–418. Providence, RI: American Mathematical Society, 1992.
- [Isham and Kuchař, 1986a] C. J. Isham and K. Kuchař. Representations of spacetime diffeomorphisms. I. Canonical parametrized field theories. *Annals of Physics*, 164: 316–333, 1986.
- [Isham and Kuchař, 1986b] C. J. Isham and K. Kuchař. Representations of spacetime diffeomorphisms. II. Canonical geometrodynamics. *Annals of Physics*, 164: 288–315, 1986.
- [Ismael, 1997] J. Ismael. Curie's principle. *Synthese*, 110: 167–190, 1997.
- [Ishibashi and Wald, 2003] A. Ishibashi and R. M. Wald. Dynamics in non-globally hyperbolic static spacetimes: II. General analysis of prescriptions for dynamics. *Classical and Quantum Gravity*, 20: 3815–3826, 2003.
- [John, 1982] F. John. *Partial Differential Equations*, 4th ed. New York: Springer-Verlag, 1982.
- [Johnson *et al.*, 2000] C. V. Johnson, A. W. Peet, and J. Polchinski. Gauge theory and the excision of repulsion singularities. *Physical Review D*, 61: 086001–11, 2000.
- [Kato, 1995] T. Kato. *Perturbation Theory for Linear Operators*. Berlin: Springer-Verlag, 1995.
- [Kay, 1992] B. S. Kay. The principle of locality and quantum field theory on (non-globally hyperbolic) curved spacetimes. *Reviews of Mathematical Physics*, Special Issue, 167–195, 1992.
- [Kay and Wald, 1991] B. Kay and R. M. Wald. Theorems on the uniqueness and thermal properties of stationary, quasifree states on spacetimes with a bifurcate Killing horizon. *Physics Reports*, 207: 49–136, 1991.
- [Koopman, 1931] B. O. Koopman. Hamiltonian systems and transformations in Hilbert space. *Proceedings of the National Academy of Sciences*, 17: 315–318, 1931.
- [Kronz, 1998] F. M. Kronz. Nonseparability and quantum chaos. *Philosophy of Science*, 65: 5–75, 1998.
- [Keller, 1957] J. B. Keller. On solutions of nonlinear wave equations. *Communications on Pure and Applied Mathematics*, 10: 523–530, 1957.
- [Lanford, 1974] O. E. Lanford. Time evolution of large classical systems. In J. Moser (ed.), *Dynamical Systems, Theory and Applications. Lecture Notes in Physics*, 38: 1–207. New York: Springer-Verlag, 1974.
- [Laplace, 1820] P. S. Laplace. *Théorie analytique des probabilités*. Paris: V. Courcier, 1820.
- [Lewis, 1973] D. Lewis. *Counterfactuals*. Cambridge, MA: Harvard University Press, 1973.
- [Lichtenberg and Lieberman, 1991] A. J. Lichtenberg and M. A. Lieberman. *Regular and Chaotic Dynamics*, 2nd ed. New York: Springer-Verlag, 1991.
- [Malament, 1984] D. Malament. Time travel in the Gödel universe. P. D. Asquith and P. Kitcher (eds.), *PSA 1984*, Vol. 2: 91–100, 1984.
- [Malament, 2004] D. Malament. On the time reversal invariance of classical electromagnetic theory. *Studies in History and Philosophy of Modern Physics*, 35B: 295–315, 2004.

- [McCall, 1995] S. McCall. Time flow, non-locality, and measurement in quantum mechanics. In S. Savitt (ed.), *Time's Arrow Today*, pages 155–172. Cambridge: Cambridge University Press, 1995.
- [Moore, 1990] C. Moore. Unpredictability and undecidability in dynamical systems. *Physical Review Letters*, 64: 2354–2357, 1990.
- [Moore, 1991] C. Moore. Generalized shifts: Unpredictability and undecidability in dynamical systems. *Nonlinearity*, 4: 199–230, 1991.
- [Nagel, 1961] E. Nagel. *The Structure of Science*. New York: Harcourt, Brace, and World, 1961.
- [Németi and David, 2005] I. Németi and D. David. Relativistic computers and the Turing barrier. *Journal of Applied. Mathematics and Computation*, forthcoming, 2005.
- [Norton, 1984] J. D. Norton. How Einstein found his field equations: 1912–1915. In D. Howard and J. Stachel (eds.), *Einstein and the History of General Relativity. Einstein Studies*, Vol. 1: 101–159. Boston: Birkhauser, 1984.
- [Norton, 1999] J. D. Norton. A quantum mechanical supertask. *Foundations of Physics*, 29: 1265–1302, 1999.
- [Olver, 1993] P. J. Olver. *Applications of Lie Groups to Differential Equations*, 2nd Ed.. New York: Springer-Verlag, 1993.
- [Payne, 1975] L. E. Payne. *Improperly Posed Problems in Partial Differential Equations*. Philadelphia, PA: SIAM, 1975.
- [Pearle, 1976] P. Pearle. Reduction of the state vector by a nonlinear Schrödinger equation. *Physical Review D*, 13: 857–868, 1976.
- [Pearle, 1989] P. Pearle. Combining stochastic dynamical state-vector reduction with spontaneous localization. *Physical Review A*, 39: 2277–2289, 1989.
- [Penrose, 1998] R. Penrose. The question of cosmic censorship. In R. Wald (ed.), *Black Holes and Relativistic Stars*, pages 103–122. Chicago, IL: University of Chicago Press, 1998.
- [Pérez Laraudogoitia, 2001] J. Pérez Laraudogoitia. Supertasks. *The Stanford Encyclopedia of Philosophy*, 2001. <http://plato.stanford.edu/archives/win2001/entries/spacetime-supertasks/>
- [Pérez Laraudogoitia, 2005] J. Pérez Laraudogoitia. Avoiding infinite masses. Preprint, 2005.
- [Pitowsky, 1996] I. Pitowsky. Laplace's demon consults an oracle: computational complexity of prediction. *Studies in History and Philosophy of Modern Physics*, 27: 161–180, 1996.
- [Planck, 1932] M. Planck. *Der Kausalbegriff in der Physik*. Leipzig: Johann Ambrosius Barth, 1932.
- [Poincaré, 1952] H. Poincaré. *Science and Method*. New York: Dover Publications, 1952.
- [Popper, 1982] K. R. Popper. *The Open Universe*. Totowa, NJ: Roman and Littlefield, 1982.
- [Pour-el and Richards, 1981] M. B. Pour-el and J. I. Richards. The wave equation with computable initial data such that its unique solution is not computable. *Advances in Mathematics*, 39: 215–239, 1981.
- [Pour-el and Richards, 1989] M. B. Pour-el and J. I. Richards. *Computability in analysis and physics*. New York: Springer-Verlag, 1989.
- [Prince, 2000] G. E. Prince. The inverse problem in the calculus of variations and its ramifications. In P. J. Vassiliou and I. G. Lisle (eds.), *Geometric Approaches of Differential Equations*, pages 171–200. Cambridge: Cambridge University Press, 2000.
- [Radicati, 1987] L. A. Radicati. Remarks on the early developments of the notion of symmetry breaking. In M. Doncel, A. Hermann, L. Michel, and A. Pais (eds), *Symmetries in physics (1600-1980)*, pages 195–207. Barcelona: Bellatera, 1987.
- [Raju, 1994] C. J. Raju. *Time: Towards a Consistent Theory*. Dordrecht: Kluwer Academic, 1994.
- [Redhead, 1987] M. Redhead. *Incompleteness, Nonlocality, and Realism*. Oxford: Clarendon Press, 1987.
- [Reed and Simon, 1975] M. Reed and B. Simon. *Methods of Modern Mathematical Physics. II: Fourier Analysis, Self-Adjointness*. New York: Academic Press, 1975.
- [Reichl, 1992] L. E. Reichl. *The Transition to Chaos in Conservative Classical Systems: Quantum Manifestations*. New York: Springer-Verlag, 1992.
- [Reichl, 2004] L. E. Reichl. *The Transition to Chaos in Conservative Classical Systems: Quantum Manifestations*, 2nd ed. New York: Springer-Verlag, 2004.
- [Rendall, 2002] A. Rendall. Theorems on existence and global dynamics for the Einstein equations. *Living Reviews in Relativity*, 2002. <http://www.livingreviews.org/Articles/Volume5/2002-6rendall>

- [Rice, 1953] H. G. Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society*, 74: 358–366, 1953.
- [Rovelli, 2004] C. Rovelli. *Quantum Gravity*. Cambridge: Cambridge University Press, 2004.
- [Saari and Xia, 1995] D. Saari and J. Xia. Off to infinity in finite time. *Notices of the American Mathematical Society*, 42: 538–546, 1995.
- [Scheffer, 1993] V. Scheffer. An invisid flow with compact support in spacetime. *Journal of Geometry and Analysis*, 3: 341–401, 1993.
- [Schnirelman, 1997] A. Schnirelman. On the nonuniqueness of weak solution of the Euler equation. *Communications on Pure and Applied Mathematics*, 50: 1261–1268, 1997.
- [Skow, 2005] B. Skow. What makes time different from space. *Noûs*, forthcoming, 2005.
- [Stachel, 1986] J. Stachel. Einstein's search for general covariance: 1912–1915. In D. Howard and J. Stachel (eds.), *Einstein and the History of General Relativity. Einstein Studies*, Vol. 1: 63–100. Boston: Birkhauser, 1986.
- [Stöltzner, 2003] M. Stöltzner. Vienna indeterminism: from Exner to Frank and von Mises. In P. Parrini, W. Salmon, and M. Salmon (eds.), *Logical empiricism: Historical and contemporary perspectives*, pages 194–229. Pittsburgh, PA: University of Pittsburgh Press, 2003.
- [Suppes, 1991] P. Suppes. Determinism or instability. Does it matter? In G. G. Brittan (ed.), *Causality, Method, and Modalities*, pages 5–22. Dordrecht: Kluwer Academic, 1991.
- [Suppes, 1993] P. Suppes. The transcendental character of determinism. In P. A. French, T. E. Uehling, and H. K. Wettstein (eds.), *Midwest Studies in Philosophy*, 18:, 242–257, 1993.
- [van Fraassen, 1989] B. C. van Fraassen. *Laws and Symmetry*. Oxford: Oxford University Press, 1989.
- [Wald, 1980a] R. M. Wald. Dynamics in non-globally hyperbolic, static spacetimes. *Journal of Mathematical Physics*, 21: 2802–2805, 1980a.
- [Wald, 1980b] R. M. Wald. Quantum gravity and time reversibility. *Physical Review D*, 21: 2742–2755, 1980b.
- [Wald, 1984] R. M. Wald. *General Relativity*. Chicago, IL: University of Chicago Press, 1984.
- [Wald, 1994] R. M. Wald. *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*. Chicago, IL: University of Chicago Press, 1994.
- [Wald, 1998] R. M. Wald. Gravitational collapse and cosmic censorship. In C. V. Vishveshwara, B. R. Iyer, and B. Bhawal (eds.), *Black Holes, Gravitational Radiation, and the Universe*. Dordrecht: Kluwer Academic Publishers, 1998.
- [Weyl, 1932] H. Weyl. *The Open World*. New Haven, CN: Yale University Press, 1932.
- [Widder, 1975] D. V. Widder. *The Heat Equation*. New York: Academic Press, 1975.
- [Xia, 1992] X. Xia. Existence of noncollision singularities in the N-body problem. *Annals of Mathematics*, 135: 411–468, 1992.
- [Zeipel, 1908] H. von Zeipel. Sur les singularités du problème des n corps. *Arkiv för Matematik, Astronomi och Fysik*, 4: 1–4, 1908.
- [Zurek, 1998] W. H. Zurek. Decoherence, chaos, quantum-classical correspondence, and the algorithmic arrow of time. *Physica Scripta*, T76: 186–198, 1998.
- [Zurek, 2003] W. H. Zurek. Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75: 715–775, 2003.

INDEX

- ABL-rule, 597
absolute objects, 1343
absolutism, 18
abundance of elements, 1221
accelerating expansion, 1209
adiabatic measurement, *see* measurement, weak
adjoint representation, 68, 69
age of the universe, 1196, 1229
Aharonov, Y., 596, 597
Albert, D., 597
Anderson, J., 1343, 1348, 1350, 1353
angular power spectrum, 1202
anomalies, 711
apparent size, 1197, 1199
area distance, 1197
Aristotle, 1313
astronomical observation, 1220, 1221
asymptotic freedom, 714
averaging scale, 1223
- background independence, 1307
baryosynthesis, 1193, 1232
BB84 (Bennett and Brassard protocol), 595, 596
BBGKY approach, 1034–1038
BCS theory of superconductivity, 493, 1122–1125
Beethoven, L., 529
beginning of the universe, 1212
Bekenstein entropy, 1294, *see* Hawking effect
Bell locality, 386, 387, 391, 395
Bell states, 567, 591
Bell's inequality, 386, 566, 596, 617, 638, 754
Bell's theorem, 308, 384, 386, 396, 755
- Bell-type experiment, 396
Belot, G., 209
Benioff, P., 556
Bennett, C., 556, 590, 593, 594, 601
Berezin function, 480
 classical limit, 480
 quantization, 459
Bergmann, P., 596, 1291
Bernoulli, D., 943
best-fit parameters, 1227
big bang, 209, 491, 1169, 1185, 1191–1194, 1209, 1210, 1212, 1213, 1214, 1236, 1287, 1318, 1322, 1424, 1425, *see* initial singularity
biorthogonal decomposition, 288, 358, 566, 578, 607, 611
bit, 558
bit commitment, 594, 600
 binding protocol, 606
 concealing protocol, 606
 'no go' theorem, 594, 606
 unconditionally secure, 601, 637, 639
black hole, 1086–1088, 1287, 1294
 thermodynamics, 1295, *see* Hawking effect
blackbody spectrum, 1077, 1197
Bloch sphere, 295, 296, 309
Bohm, D., 323, 385, 418, 437, 566, 653
Bohm interpretation of quantum mechanics, 373–377, 392, 646, 651–653, 762, 1404, 1406, *see* hidden variables
Bohr, N., 424, 433–444, 555, 647
Bohr's atomic model, 424, 483
Bohr's correspondence principle, 426,

- 429
- Boltzmann, L., 952–992
- Boltzmann equation, 962–969, *see*
 BBGKY approach
 Lanford’s approach, 1028–1033
- Boltzmann’s combinatorial argument
 and responses, 975–983, 983–
 992
- Bondi, H., 650
- Boolean algebra, 642
- Born rule, 524
- Bose–Einstein
 condensation, 1082–1084
 statistics, 848–855
- Bose gas, 1125
- bra-ket notation, *see* Dirac bra-ket
 notation
- braided tensor category (BTC*), 870
- braiding, 867
- Brassard, G., 556, 590, 594, 601, 650
- Briegel, H., 616
- broadcasting, 580
 ‘no broadcasting’ theorem, 581
- Broglie-Bohm theory, *see* Bohm in-
 terpretation of quantum me-
 chanics
- Bronstein, M., 1291
- Brout–Englert–Higgs mechanism, 693–
 698
- Brukner, Č., 632
- BTC* (braided tensor category), 870
- Bub, J., 632, 646
- Bub–Clifton theorem, 646
- Butterfield, J., 138, 155, 209, 397
- C^* -algebra, 419, 445, 451–463, 468–
 474, 492–514, 520–522, 633–
 637, 651, 653, 732–740, 751,
 754, 755, 841, 843–845, 894,
 1110–1113, 1120, 1144, 1151,
 1163, *see* GNS construction,
 von Neumann algebra
- canonical action, 103
- canonical quantization, 447
- Carnot’s theorem, 934
- category, 863
 *-category, 869
 C^* -, 874
 \mathbb{C} -linear, 869
 Ab-, 868
 Abelian, 875
 additive, 869
 braided monoidal, 867
 equivalence, 864
 complex-linear, 869
 monoidal, 865
 of localized transportable endo-
 morphisms, 783f
 product, 865
 symmetric monoidal, 867
- Cauchy
 development, 1415–1418
 horizon, 1416–1417
 surface, 1394, 1409–1411, *see* global
 hyperbolicity
- causal properties of relativistic space-
 times, 263
 past distinguishing, 271
 chronology, 271
 closed timelike curves, 271, 1421
 future distinguishing, 271
 global hyperbolicity, 1410, 1419,
 1423–1424, *see* Cauchy sur-
 face
 past distinguishing, 271
 stable causality, 1368, *see* global
 time function
 strong causality, 271, 1410
- CBH (Clifton–Bub–Halvorson) theo-
 rem, 632, 650
- CBR (cosmic background radiation),
 1192, 1195, 1197, 1200, 1202–
 1206, 1215, 1223, 1226, 1231
 anisotropies, 1209, 1214, 1223, 1224,
 1227, 1231
 isotropy, 1195, 1225, 1226
 polarization, 1210
 power spectrum, 1202, 1203, 1217

- temperature, 1229
- CCR (canonical commutation relations), 287, 334, 340–341, 447, 1125, 1423
- Chalmers, A.,
 - on Curie’s principle, 1334
- channel capacity, 562, 564
- chaos
 - classical, 490, 491, 520, 1390, 1403
 - deterministic, 1390, 1391
 - quantum, 478, 486–491, 519, 1106, 1402
 - intrinsic, 490
- Chau, H., 594, 601
- chronology, *see* causal properties of relativistic spacetimes
- chronology protection conjecture, 1419
- Church, A., 614
- Church–Turing thesis, 614, 1420, *see* Turing machine
- classical limit (of quantum mechanics), 360, 471–530, *see* high temperature limit, thermodynamic limit, van Hove limit
 - the limit $\hbar \rightarrow 0$, 471–491, 1156–1157
 - the limit $N \rightarrow 0$, 492–514
- classical observables, 493, 512, 514ff, 1162–1163
- classical probabilities, 283, 285, 309
- classical states, 316, 422, 472, 500, 514
- classical world, 420
 - absolute existence of, 420
- Clausius, R., 932–938, 944
- Clifton, R., 632, 646
- clock (or twin) paradox, 236
- cloning, 582
 - ‘no cloning’ theorem, 580, 581
 - and measurement, 652
- closed timelike curves, 271, 1419 *see* causal properties of relativistic spacetimes 1419
- clustering, 740, 1108, 1109
- CNOT gate, 616, 632
- co-adjoint representation, 73
 - symplectic structure of, 99
- coarse-graining (for consistent histories), 526
- coherent state, 429, 455, 472, 519
 - Schrödinger’s, 455
- cokernel, 875
- collapse interpretations (of quantum measurement), 377–382, 645–647
- collapse postulate, 363, 377, 378, 397, 1165
- common cause, 382–386
- complementarity, 441, 555, 647
 - early Bohr, 441
 - Heisenberg, 442
 - Landsman, 443
 - later Bohr, 442
 - Pauli, 443
- completely positive map, 456, 574, 755
- completeness condition, 570
- complexity class
 - EXP**, 614
 - NP**, 615
 - P**, 614
- compression/decompression scheme, 561, 589
- computational complexity, 614
 - easy (tractable), 614
 - hard (intractable), 614, *see* complexity class
- concordance model, 1214–1216
- Connes, A., 1295
- confinement, 721
- conservation of energy-momentum (in general relativity), 240–243
- conservation principles (in general relativity), 252–256, *see* Noether’s theorems
- consistent histories, 419, 422, 517, 518, 522–529
- constant mean curvature time, 214–

- 221
- constraint equation, 1230, 1352, 1372, 1412, 1424
- constraints (in Hamiltonian mechanics), 151–153, 180
- first class, 180, 1379
 - primary, 1379
 - secondary, 1379
- contact transformations, 1341
- container view of spacetime, 18, 1312, 1377, 1411
- continuous field of C^* -algebras, 457
- continuous spontaneous localisation, 380
- Copenhagen interpretation, 349, 422, 433–446, 555, 645, 655, 1257, 1232
- Copernican Principle, 1225, 1226, 1273
- correspondence principle, 418, 424–429, 483
- cosmic censorship conjecture, 1419, 1425
- cosmic variance, 1203, 1217
- cosmological constant, 243–249, 1187, 1188, 1190, 1197, 1200, 1202, 1227, 1408, *see* Einstein's field equations
- cosmological principle, 1223, 1231
- cosmological time function, 216
- cotangent bundle
- symplectic structure of, 9, 122
- cotangent lift, 18, 20, 58, 116
- covariance, 330, 335, 1343
- general, *see* general covariance
- Crépeau, C., 590
- creation of the universe, 1212, 1218, 1230, 1236, 1271
- critical density, 1188, 1199, 1201
- Curie's principle, 1333–1336, 1346, 1374–1375
- Curie, P., 1374–1375, 1415
- cyclic (ignorable) coordinates, 5, 15
- Darboux's theorem, 10, 85, 91, 97, 98, 142
- dark energy, 1185, 1188, 1200, 1202, 1216, 1224, 1227, 1232, 1267, 1268, 1408
- dark matter, 1201
- cold dark matter (CDM), 1187, 1189, 1201, 1202, 1244
 - hot dark matter, 1201
- de Broglie-Bohm theory, *see* Bohm interpretation of quantum mechanics
- De Finetti's representation theorem, 500
- de Sitter spacetime, 213, 215, 1213,
- de Sitter universe, 1189, 1190, 1198, 1204, 1212
- Debye, P., 1080
- deceleration parameter, 1187, 1197, 1198, 1227
- decoherence, 211, 365–368, 372, 373, 515–529, 646, 654, 1401, 1405
- deformation quantization, 456
- functorial reinterpretation, 470
- delay-differential equations, 1371
- Deligne, P., 901, 919
- density operator, 281–285, 288, 289, 295, 296, 302, 307, 313, 317, 376
- density parameter, 1187, 1188, 1227
- Descartes, R., 1312–1313
- Deser, S., 1292
- determinism
- as a methodological imperative, 1372
 - varieties of, 1373–1374
 - Laplacian, 1370
 - non-Laplacian, 1373–1374
- Deutsch, D., 556, 617, 631
- Deutsch's XOR algorithm, 617, 618
- Deutsch–Jozsa algorithm, 618, 621
- Dieks, D., 580
- diffeomorphism, 31
- diffeomorphism invariance, 1308, *see* general covariance

- differential equation
 - ordinary, 32, 1373, 1385
 - partial
 - elliptic, 1382
 - hyperbolic, 1395, 1396
 - parabolic, 1382
 - weak solutions, 1384
- dimension in a category, 871
 - integrality of, 907
- Dirac bra-ket notation, 296–198, 305
- Dirac equation, 677–679, 720
- Dirac matrices, 469, 678, 687
- Dirac operator, 471, 1295, 1296
- Dirac picture, 306, 1161
- Dirac, P. A. M., 153, 377, 378, 427, 429, 432, 447, 459, 462, 467, 469, 470, 473, 555, 642, 718, 1239, 1291, 1379
- disjoint representations, 498, 732, 843, 844
- disjoint states, 498, 843
- distances of observed objects, 1195
- domain of causal influence, 1205
- domain of dependence, 241, 1387m
 - see* Cauchy development
- domain of prediction, 1397–1398, 1419
- Doplicher, S., 744, 766, 800, 801, 815, 837, 840, 845, 855
- Doplicher–Haag–Roberts (DHR)
 - theory of superselection, 502, 781ff
- double slit experiment, 345, 377, 441, 443, 524, 536
- dual of a Lie algebra, 88, 90
- dynamical objects, 1343
- dynamical reduction theories, 379
- dynamical triangulation, 1288, 1296
- Earman, J., 177, 197, 206, 209
 - on Curie’s principle, 1335
- Eddington, A. S., 514, 1088
- EGS theorem, 1226
- Ehrenfest time, 478
- Ehrenfest’s theorem, 475
- Ehrenfests, on Boltzmann, 930, 980, 1005
- eigenstate–eigenvalue link, 285, 286, 357–359, 367, 378, 380, 381, 1405–1406
- einselection, 516–522, *see* decoherence
- Einstein, A., 206, 252, 258, 259, 348, 350, 418–424, 427, 433, 435, 442, 471, 478, 486, 488, 555–557, 639, 640, 645, 647–652, 940, 941, 977, 989, 991, 993, 1003, 1021, 1075–1082, 1290, 1309–1314, 1346–1353, 1356, 1357, 1363, 1410–1414
- Einstein Field Equations (EFE), 200, 243–245, 249, 263, 1185, 1187, 1190, 1224, 1230, 1290, 1352, 1409–1416
- Einstein’s ‘hole argument’, 1308, 1349, 1351, 1411
- Einstein’s principle of relativity, 1346
- Einstein static universe, 1184, 1189, 1212
- Einstein–de Sitter model (or universe), 1189, 1198, 1204
- Ekert, A., 596, 599
- embedding, 41
- empirical content of quantum theory, 307–308, 321–339
- energy-density conservation equation, 1186
- energy conditions (in general relativity), 240, 1207
 - dominant, 240, 1394
 - strong, 1187
- ensemble interpretations,
 - of quantum mechanics, 363, 432, 478
 - of statistical mechanics, 943
- entangled state (in quantum mechanics), 288, 367, 368, 396, 445, 565, 566, 754–756
- entanglement (quantum), 397, 439, 444–446, 518, 555, 556, 590–

- 594, 602, 603, 605, 611, 612, 616, 630, 635–640, 647, 649–651, 754–756
- entropy
 conditional, 562, 563, 575, 577
 gravitational, 1208
 joint, 563, 577, 578
 relative, 564
 Shannon, 557, 558, 562, 564, 576, 577
 von Neumann, 575, 577, 579, 1091
- environment-induced superselection, 516–522, *see* decoherence
- EPR (Einstein–Podolsky–Rosen) argument (or paradox), 308, 349–355, 385, 555, 556, 590, 639
 EPR criterion for physical reality, 351–353, 384
 EPR attack (or cheating strategy), 601, 603, 605, 606, 609, 610, 637
- equivariance, 63, 68, 75
 of momentum maps, 103, 112–114, 117
- ergodic theory (classical) 1005–1020, 1025, *see* quantum ergodic theory
 ergodic hierarchy, 1011–1020, 1104–1105, 1390
 ergodic hypothesis in early Boltzmann, 955–962, *see* Boltzmann’s combinatorial argument *and* H-theorem
 ergodic theorem, 1006–1009, 1093
 Khinchin’s approach, 1017–1028
- Euclidean group, 17–20, 62, 66, 103, 1100, 1113, 1151, 1154
- Euler
 equations for rigid body, 81, 91
 equations for fluid, 16, 1383–1384
- Euler–Lagrange equations, 145, 146, 158, 159, 184, 190, 208, 688, 1356–1358, 1378, 1379
- Everett interpretation, 370, 646, 1407–1408
 Saunders-Wallace version, 646
- evolution of the universe, 1185, 1188, 1233, 1267, 1273, 1319
- existential interpretation (of quantum mechanics), 517
- expansion of the universe, 1183, 1184, 1187, 1192, 1196, 1200, 1214, 1221, 1239, 1245, 1252,
- Faddeev, L., 689–691, 695, 1292
- Faraday lines, 1300
- Faraday, M., 1300, 1314, 1324
- Fermi–Dirac statistics, 679, 704, 719, 845–855, 1085
- Feynman, R., 556, 1170, 1291, 1292, 1391
- Feynman diagram, 664–667, 676, 682, 699, 704, 709
- Feynman path integral, 671–674
- Feynman propagator, 666, 698
- Feynman rules
 classical theory, 663–669
 quantized theory, 674–677
 Dirac fields, 681–683
 calculation of scattering matrix, 690–691
 perturbive quantum general relativity 1292
- fiber functor, 881
 existence, 884, 900, 902
 uniqueness, 882, 893
- fidelity, 561, 588, 589
- Fierz, M., 1290
- fine-graining (for consistent histories), 526
- fine tuning, 723, 1205, 1209, 1212, 1215, 1218, 1238, 1248, 1249, 1264, 1265, 1285
- finitely generated
 algebra, 918
 tensor category, 900
- folium, 498

- Friedman, M., 197, 206
 Friedmann equation, 1187, 1188, 1196, 1198, 1211, 1217
 Friedmann–Lemaître (FL) model (or universe), 1184, 1187, 1188, 1190, 1195, 1204, 1207–1209, 1212, 1215, 1222–1229, 1234, 1238, 1266, 1269, 1270, 1275
 Friedmann–Robertson–Walker model, 1187, 1322, 1426, 1427, *see* Friedmann–Lemaître model
 Frobenius’ theorem, 7, 29, 38–43, 51, 94, 97, 98, 151, 250, *see* Perron–Frobenius’ theorem
 Fuchs, C., 580, 601, 632, 650
 functor, 864
 - *-preserving, 869
 - braided monoidal, 867
 - essentially surjective, 864
 - faithful, 864
 - fiber, 881
 - full, 864
 - monoidal, 866
 - symmetric monoidal, 867
 funnel property, 743
 future distinguishing, *see* causal properties of relativistic spacetimes

 Gödel, K., 614
 Galilean group, 25, 299, 329–334
 Galilean transformation, 299, 337, 649, 1375, 1397, 1398
 Galois, É., 1337–1338
 gauge conditions
 - Lorentz gauge, 1378, 1379
 gauge fixing, 153, 688, 692, 704, 716, 1318, 1378
 gauge freedom, 151–153, 177–191, 1352, 1376, 1378–1380, 1410, 1413
 gauge symmetry, 197, 203, 467, 685, 687, 719, 806, 809, 811, 815, 830, 846–848, 853, 854, 1124, 1355, 1359, 1378–1380, 1410, 1413
 general covariance, 133, 136, 196–219, 1302, 1306, 1310, 1311, 1343, 1348–54, 1357, 1358, 1363, 1409, 1411, 1414
 general linear group, 44, 49, 409
 geometric quantization, 462–467, 471, *see* quantization
 - functorial reinterpretation, 470
 geometry of the universe, 1222
 Ghirardi, G.-C., 379–381, 420, 522, 645, 1403
 GHZ (Greenberger–Horne–Zeilinger) theorem, 361, 387
 Gibbs’ statistical mechanics, 992–1005, *see* BBGKY approach
 Gisin, N., 379, 569
 Gleason’s theorem, 281, 282, 286, 295, 308, 634, 644, 654, 1089, 1405
 global hyperbolicity, 1410, 1421–1424, *see* Cauchy surface
 global observable, 1101, 1109, *see* macroscopic observables
 global time function, 1370, 1410, 1411, 1419
 GNS (Gelfand–Naimark–Segal) construction (or representation), 318, 442, 445, 498, 499, 502, 503, 505, 508, 510, 732, 740, 758, 759, 782, 800, 806, 827, 829, 830–833, 836, 843, 845, 851, 1089, 1101, 1109, 1110–1113, 1117, 1119, 1120, 1121, 1124–1127, 1133, 1144, 1148, 1152, 1154, 1163
 Goroff, M., 1292
 gravitational lensing, 1197–1199, 1202, 1203, 1214, 1227
 gravitational instability, 1193, 1210, 1214
 gravitational waves, 1186, 1203, 1207, 1210, 1232, 1252, 1290, 1291
 GR, *see* GTR (general theory of relativity)

- Green, M., 1288, 1298
 groups and group theory, 299, 409, 685, 1337–1342
 amenable, 740, 1113
 continuous, 44, 1340–1342
 discrete, 1340
 Lie, 7, 8, 15, 19, 20, 29, 37, 40, 44–79, 147, 1341, 1342, 1376
 matrix, 44, 50, 51, 71, 72
 semi-simple, 1113
 compact, 900, 901
- GRW (Ghirardi–Rimini–Webber) ‘collapse theory’, 645, 651
- GTR (general theory of relativity), 133–138, 147, 156, 172, 178, 184, 196–221, 229, 246, 263, 384, 467, 724, 1185, 1230, 1287, 1348–1354, 1409–1423, *see* Einstein field equations
- Gupta, S., 1290
- Gutzwiller trace formula, 487, 488
- H*-theorem, 962–969
 objections and responses, 970–974
- Haag, R., 738ff, 777, 781ff, 1099, 1124, 1125
- Hadamard gate, 616, 632
- Hadamard transformation, 618, 620–622, 624–626
- Hadamard, J., 1105, 1390
- Halvorson, H., 569, 632, 633
- Hamilton’s equations, 11, 85, 87, 90, 141
- Hamilton–Jacobi equations of GR, 1293
- Hamiltonian approach to mechanics and field theory, 8–13, 79–119, 140–143, 147–154, 165–168, 180, 193
- Hamiltonian, 303, 304–307, 323, 325
- Hamiltonian symmetries, 12, 58
- Hamiltonian system, 12, 150
- Hamiltonian vector field, 12, 85, 458
- Hardy’s axioms (for quantum mechanics), 320–322
- Hardy’s no-go theorem, 387–389
- Hartle, J., 518, 522, 528, 530, 546, 1262, 1264, 1294, 1315, 1323
- Hartle–Hawking ‘no-boundary’ proposal, 1213, 1236
- Hausdorff property (of topological spaces), 63, 147, 155, 175, 445, 457, 507, 635, 885, 894, 1295, 1409, 1418
- Hawking effect, 1143, 1144, 1294, 1424, *see* Unruh effect
- Hawking, S., 244, 271, 1143, 1207, 1213, 1236, 1294, 1424
- Healey, R., 191, 209, 375, 421
- heat equation, 927, 1096, 1380–1381
- Heisenberg, W., 340, 342, 417, 418, 420, 421, 426–448, 459, 471, 484, 486, 515, 523, 529, 555, 641, 1126, 1291, 1292
- Heisenberg cut, 418, 437–441, 446, 518
- Heisenberg picture, 306, 307, 448, 475, 520, 521, 523, 635, 1115, *see* Schrödinger picture
- Heisenberg equation, 307, 509
- Heisenberg’s matrix mechanics, *see* matrix mechanics
- Helmholtz, H. von, 159, 339, 1381
- Helmholtz free energy, 939, 1091
- Henneaux, M., 153, 467, 1379–1380
- Hepp, K., 477, 492, 493, 499, 502, 512, 513, 1163
- Hertz, H., 28
- Hertz, P., 995
- hidden variables (in quantum mechanics), 372–376, 384–395, 431, 446, 646, 648, 653, 754, 1401–1406, *see* Bohm interpretation
- Higgs mechanism, *see* Brout–Englert–Higgs mechanism
- high temperature limit, 1078, 1081,

- 1157–1159
- Hilbert, D., 1076, 1268, 1270, 1348, 1351, 1357, 1361, 1363
- Hilbert space, 397–400
 formalism for quantum statistical physics, 1089–1094
 limitations of, 1094–1099
- hole argument, *see* Einstein's hole argument
- Holevo bound, 582, 585, 593
 derivation, 583
- holism, 395
- horizons, *see* physics horizon
 event, 1203, 1417
 Cauchy, 216, 1414, 1415
 causal, 1203, 1204, 1209
 particle, 1203, 1204, 1206, 1209, 1215, 1230
 visual, 1203, 1204, 1206, 1209, 1226, 1227, 1230, 1231, 1247, 1272
- horizon problem, 1204, 1209, 1238
- hot big bang, *see* big bang
- Hubble constant (or parameter), 1186, 1187, 1190, 1196, 1197, 1227
- Hughston–Jozsa–Wootters theorem, 569, 590
- Husimi function, 480
- hydrodynamics as part of statistical physics, 926
- Hyperion, 418, 478, 519, 1403
- ideal fluid, 8, 15, 16, 82
- ignorable coordinates, *see* cyclic coordinates
- ignorance interpretation (of mixed states), 357, 432, 433, 439, 440, 512, 516, 544, 749
- Immirzi parameter, 1295, 1301
- improper mixture, *see* mixed state, improper
- impulsive measurement, 321–324, 327, 337, 377
- incompatible observables, 285, 308, 346–349
- incomplete motion, 491, 492, 1382, 1385, 1400
- inductive limit C^* -algebra, 496, 504, 739, 765, 829, 879, 900, 901, 911, 920, 921, 1099, 1100, 1107, 1129, 1141, 1151, 1152
- inertial frame, 140, 164, 169, 170, 188, 206, 334–339, 350, 451, 648, 649, 1297, 1375, 1376
- inflationary model (or universe), 1209–1212, 1215, 1224, 1233, 1234, 1284
 chaotic, 1206, 1210, 1213, 1219, 1227, 1231, 1234–1236, 1244, 1245, 1256, 1262, 1266, 1269
- information
 accessible, 582
 channel, 562
 classical, 557, 576, 577
 compressibility, 576, 577
 compression, 562, 578
 'block coding', 561
 blind, 586
 classical, 557
 quantum, 585
 visible, 586
 compression rate, 560, 586
 inaccessible, 577
 mutual, 562, 564, 575, 577
 quantum, 565, 576, 577
 represented physically, 576
 Shannon's sense, 575
 source, 557, 561
 ergodic, 558
 stationary, 557
- information loss paradox, 1424, 1426
- initial singularity (of the universe), 1188, 1190, 1204, 1211, 1232, 1235, 1236, 1284, 1426, *see* big bang
- initial conditions for the universe, 1223, 1233, 1237, 1238, 1273, 1426
- initial value formulation/problem, 1370,

- 1381, 1385
 properly posed, 1390
 well posed, 1388–1390, 1395, 1414
 instantons, 719
 integrable system, 482, 483, 485
 classical limit, 482
 integral manifold, 42, 43, 98, 151
 interaction Hamiltonian, 322, 324, 336, 364
 interaction picture, 306, 1161
 interference pattern, 345–347, 378, 441,
 see wave-particle duality
 intermediate isotropisation, 1226
 invariance, 328, 333, 337, 1331, 1343
 invariant means and states, 740, 1108
 involution, 42
 Isham, C., 209, 452, 462, 522–525,
 1212, 1254, 1257, 1294, 1313,
 1318, 1324
 isotropic universe, 1225, 1226, *see*
 Copernican Principle
 isotropy group, 65, 66, 68, 120, 128

 Jacobson, T., 1293, 1303
 Jaynes, E., 569, 932, 943, 977, 993,
 1003, 1010, 1090
 Jozsa, R., 557, 569, 580, 585, 588,
 590, 616, 618, 621, 627, 630–
 632

 Kaluza–Klein theory, 28, 29
 Kent, A., 517, 522, 527, 602, 608
 Kepler, J., 1305
 Kepler problem, 5, 15, 174–175, 483,
 484
 key distribution and storage, 556, 594,
 595–596, 599
 public, 617
 secure, 602, 650
 Killing fields, 165, 199, 215, 252–256,
 1421
 kinematic independence, 633, 635, 636,
 639, *see* microcausality
 kinematics on Lie groups, 75
 kinetic theory of gases, 927, 941–952

 Klein, F.
 and group theory, 1338–1342
 and his Erlangen Program, 1339–
 1340
 Klein, M., 649, 930, 931, 957, 962,
 965, 974, 990
 Klein, O., 1290
 Klein-Gordon equation, 154, 192, 207,
 208, 1394, 1396, 1416, 1422
 KMS (Kubo–Martin–Schwinger) states,
 1114–1154
 extremal KMS states and pure
 phases, 499, 1144–1154
 in BCS theory, Bose gas and Heisen-
 berg model, 1122–1128
 in QFT, 1142–1144
 stability conditions, 1128–1142
 Tomita–Takesaki theory, 1119–
 1122
 toy-model, 1115–1119
 Kochen–Specker theorem, 359–362, 387,
 1405
 Kolmogorov axioms, 644, 1089
 Kosso, P., 1350
 Kraus operator, 570, 756
 Kretschmann, E., 206, 1353
 Kuhn, T., 423

 Lagrange, J. L., 1337, 1345
 Lagrangian approach to mechanics and
 field theory, 143–146, 154–
 164, 179, 192
 Lagrangian submanifold, 485
 Landau, L., 417
 Landsman’s theorem, 320
 Lanford’s theorem, 1031–1033
 Laplace’s demon, 1388, 1392
 Laplace, P.-S., 942–943, 1388
 large scale of the universe, 1220
 Last Scattering Surface, *see* LSS
 lattice of propositions, 309–312, 403–
 405, 641
 laws of creation of the universe, 1218
 laws of nature

- empiricist conception, 1371
- local algebras, 745–750, *see* global observables
- Lebowitz, J., 596
- left action, 56
- left translation, 45
- left-invariant function, 123
- Leggett’s argument against decoherence, 516
- Leibniz, G., 1313
- Lewis, D., 155, 1371
- Lie algebra, 35, 1341
- Lie algebra of a Lie group, 45, 46
- Lie bracket, 37
- Lie derivative, 33
- Lie group, *see* groups and group theory, Lie algebra
- Lie groupoids, 462
- Lie subgroup, 50
- Lie, S., 15
 - and group theory, 1338–1342
- Lie–Poisson bracket, 88, 124, 125
- Lie–Poisson reduction theorem, 19, 77, 82, 88
- limit $\hbar \rightarrow 0$, *see* classical limit
- limit $N \rightarrow \infty$, *see* classical limit
- limit to size of structure, 1205
- Linden, N., 631
- Lipschitz condition, 1383, 1387
- Lo, H.-K., 594, 601
- locality, 308, 351, 381, 391, 396
- localizability of particles, 481
- Lorentz theory, 649, 651
- Lorentz, H., 417
- Lorentz-invariance, 383, 388, 396
- LSS (Last Scattering Surface), 1192, 1202, 1206, 1214, 1230
- Lüder’s rule, 284, 289, *see* collapse postulate

- Mackey’s theorem, 331–334, 447–452, 461–462, *see* system of imprimitivity
- Mackey, G., 447

- macroscopic average, 494
- macroscopic observables, 493–496, 504–506, 512, *see* global observable
- magnetic monopoles, 718
- magnitude-redshift relation, 1214
- Malmquist bias, 1199
- many universes, 1219, 1259–1266
- many-somethings interpretation of quantum mechanics, 370, 646, 1407–1408
- Marsden–Weinstein reduction, 467
- Marsden–Weinstein–Meyer theorem, 127
- matrix mechanics, 277, 417, 427, 428, 641
- matrix representation of quantum states, 282, 298, 364
- Maudlin, T., 183, 184, 191, 209
- Maxwell, J., 944–952, 1300
- Maxwell’s equations, 188–190, 1395, 1416
- Mayers, D., 594, 601, 608
- measurement, 321–325, 327, 328, 334–337, 345, 352, 355–359, 362, 569
 - first kind, 324
 - generalized, 569, 573, 579
 - impulsive, 323, 326, 338, 376
 - non-selective measurement, 394, 575
 - operator, 323, 570–572, 579
 - projective, 570, 571
 - and entropy, 579
 - repeated, 311, 325
 - second kind, 324
 - selective, 575
 - weak, 324–327
- measurement problem, 275, 355–368, 376–381, 431–433, 512, 518, 526, 1163–1168, *see* Schrödinger’s problem
- microcausality, 739, 741–743, 746, 749–753, 766, 771, 772, 807, 839,

- see* kinematic independence
- microlocal analysis, 419, 476, 482
- microscopic quantum observables, 506
- Milne model, 1189
- minimum apparent size, 1198
- Minkowski spacetime, 256–263
- Minkowski, H., 1363
- misconceptions about cosmology, 1214
- Misner, C., 1291
- mixed state, 284, 289, 290, 316, 356, 364, 430–433, 439, 473, 479, 494, 501, 502, 565, 568–571, 574, 576–578, 588–592, 636, 639, 643, 749, *see* ignorance interpretation
 - improper, 289, 366
 - maximally mixed, 749
 - proper, 289, 366, 605
- mixture, *see* mixed state
- modal interpretations of quantum theory, 375, 646, 761–766, 1406–1407
- model selection, 1229
- module, 876
- momentum(or non-relativistic quantum mechanics), 287, 329, 335, 340, 349–352
- momentum function, 114
- momentum map, 6, 106, 464
- momentum map on cotangent bundle, 114
- monic, 874
- monoid, 876
 - absorbing, 902, 913
 - module, 876
 - quotient, 878
- n -Body problem, 175–177, *see* Newtonian n -body problem
- Naimark’s dilation theorem, 453
- naive realism, 359
- natural transformation, 864
 - monoidal, 866
- Navier–Stokes equations, 1383–1384, 1396
- Neumann, J. von, 430, 432
- Newton, I., 1312, 1314
- Newton’s equations of motion, 1385
- Newton–Wigner position operator, 451
- Newtonian n -body problem, 139–146, 1385
- Newtonian gravitation theory, 245, 263–270, 1385
- Nieuwenhuizen, P. van, 1292
- ‘no go’ theorems, 383, 645
- no-signaling theorem, 391
- Noether’s theorems, 4, 6, 11, 108, 161, 253, 1355–1359, 1378, 1411
- noisy channel, 562, 564
- non-baryonic matter, *see* dark matter
- non-contextuality, 360, 387
- non-locality, 308, 381–397
 - in quantum field theory, 750–756
- non-localizability of particles, *see* localizability of particles
- non-singular start of universe, 1211
- Norton, J., 197
- NP**-complete, 615
- nucleosynthesis, 1192, 1199, 1214, 1222, 1232
- null cone observation, 1194
- number count dipole, 1230
- number counts, 1198, 1199, 1214, 1224, 1227
- object
 - conjugate, 869
 - irreducible, 870
 - projective, 875
- observable, 277, 282, 283, 289, 312, 328, 337, 359, *see* POVM, PVM
 - positive-operator-valued measurement, 277, 569–575
- observable universe, 1183, 1198, 1206, 1222, 1230
- observational cosmology theorem, 1222

- observational completeness, 1232
- observed isotropy, 1223
- Ohanian, H., 1343
- one-time pad, 594
- operators introduced, 277f., 401–403
 - at a point in quantum field theory, 766–777
- origin of the light elements, 1193
- origin of the universe, 1212, *see* big bang
- outcome independence, 390, 395

- Painlevé’s conjecture, 1385
- parameter independence, 390–395
- partial observable, 1308, 1317
- particles in quantum field theory, 756–761, *see* Unruh effect
- passage to the quotients, 63, 76
- past distinguishing, *see* causal properties of relativistic spacetimes
- Pauli, W., 443, 555, 1290
- Pearle, P., 379, 381, 645, 1403
- Penrose, R., 1419
- Peres, A., 590, 608, 1293
- perfect (anti-)correlation, 288, 323, 349, 365, 381, 383
- permutation-invariant state, 500
- Perron–Frobenius theorem, 1048, 1050
- phase gate, 616, 632
- phase space quantization, 452, 456
- phenomenon for Bohr, 434, 442
- Phoenix universe, 1213
- physical proposition, 310, 311–312, 316–317
- physics horizon, 1232, 1233
- $\pi/8$ gate, 616, 632
- Piron’s theorem, 312
- Planck scale, 1287, 1289
- Planck, M., 423, 1075–1079, 1370
- pocket universe, 1210, *see* inflationary model, chaotic
- Podolsky, B., 349

- Poincaré’s recurrence theorem, 983–984
- pointer state, 323, 519
- Poisson antimorphism, 104, 105, 108, 115
- Poisson bracket, 84, 307, 318, 319
- Poisson equation, 1379
- Poisson manifold, 6, 14, 24, 79, 84, 317, 458
 - quotient of, 101
- Poisson map, 93
- Poisson reduction theorem, 101
- Poisson space, 506
- Poisson submanifold, 95
- Poisson structure, 91
- polarization spectrum, 1203
- Polyakov, V., 1300
- Popov, V., 1292
- position in non-relativistic quantum mechanics, 330–333, 335–338, 340, 349–352, 375, 382, 451
- POVM (positive operator valued measure), 277–283, 321, 327, 453, 569–575
- pre- and post-selected states, 326, 596
- predictability, 1388, *see* domain of prediction
 - and chaos, 1388–1389
 - and determinism, 1386–1390
 - in general relativistic physics, 1419
 - in quantum mechanics, 1402–1403
 - in special relativistic physics, 1396–1398
- preferred basis, 372
- prequantization, 463
 - line bundle, 464
- presymplectic manifold, 151, 172, 179–180, 189
- primary states, 498
- Primas, H., 444
- Primas–Zurek cut, 518
- prime factorization, 557
- principle vs. constructive theories, 647–652

- probabilities in the Copenhagen interpretation, 440
- probability, 279, 282, 283, 298, 300, 318, 929, 941–943
- problem of time, 133, 136, 137, 196–221, 1316
- projective representation, 301
- projector morphism, 869
- proper map, 64
- protective measurement, *see* measurement, weak
- public-key cryptography, 617
- pure state in quantum mechanics, 284, 285, 289, 293, 296, 301, 302, 312, 314–317, 320–322, 358
- purification, 568, 569, 574
- PVM (projection-valued measure), 278, 279, 281, 331, 571

- quantization, 153, 154, 177, 195, 220–221, 304, 446, 467, 669, *see* geometric quantization *and* system of imprimitivity *and* Weyl quantization
- quantum information and computation, *see* key distribution, bit commitment, information algorithms, 617
 - communication, 556, 590, 592
 - computation, 295, 556, 614
 - computer, 556
 - circuit model, 616
 - cluster state, 616
 - one-way, 616
 - cryptography, 556, 594, 595
 - dense coding, 593
 - information, 565, 576, 577, *see* information
 - operation, 574
 - parallelism, 631
 - qubit, 577
 - teleportation, 556, 590, 591
- quantum cosmology, 1183, 1212, 1294
- quantum ergodic theory, 1092–1094
 - for time, 1102–1106
 - for space, 1106–1114
- quantum gravity, 1211, 1212, 1214, 1287
- quantum logic, 309, 369
- quantum object in Primas' sense, 444
- quantum postulate, 438
- quantum probability theory, 283, 1089
- quantum theory of infinite systems, 492, 496, 738–750, 1094–1102
- quantum world, 434, 526
- quasi-equivalent representations, 498, 732, 745, 843, *see* representation
- quasilocal observables, 496
- quintessence, 1200
- quotient space, 6, 62, 63

- Raggio's theorem, 445
- randomized algorithm, 629
- Raussendorff, R., 616
- Raychaudhuri equation, 1186, 1188, 1190, 1211
- Redhead, M., 191
- redshift-distance relation, 1184, 1196–1198
- reduced density operator, 288, 366, 367, 375, 568
- reduced phase space, 14, 21, 24, 128, 152, 183, 189, 202
- reduced state, *see* reduced density operator
- reduction of dynamics, 102, 126, 129
- Reeh–Schlieder theorem, 742–743
- reference frame, 329, 335, 337, 339, 347
- Regge calculus, 1296
- relationism, 17–18, 1312
- relative configuration space, 18, 21
- relative phase space, 21
- relative states, 370, 568
- relativistic spacetime structure
 - causal structure, *see* causal properties of relativistic space-

- times
 - conformal structure, 232–233, 270–271
 - projective structure, 232–233
- reliable distance measurement, 1221
- remote steering, 590, 591, 637, *see* biorthogonal decomposition
- renormalization, 683, 705–711, 714
- representation, 57, 409
 - inequivalent representations, 778f.
- right action, 56
- right translation, 45
- right-invariant function, 123
- right-invariant vector field, 47
- rigid body, 15, 61, 80, 90
- Rimini, A., 645
- Roberts, J., 800, 801, 815, 837, 840, 845, 855, 886
- Robertson-Walker cosmological models, 1186–1188, 1207, 1215, 1223–1227, 1231
- robust state, 518
- Rosen, N. 349
- Rosenfeld, L., 1290
- rotation in relativity theory, 249–252
- rotation group, 44, 52, 60, 72, 74, 90, 98
- RSA, 617
- Ruetsche, L., 221
- Ruffini, R., 1343
- Rynasiewicz, R., 206

- Sagnotti, A., 1292
- Saunders, S., 646
- scars, 490
- Schmidt decomposition, *see* biorthogonal decomposition
- Schrödinger, E., 418, 428, 556, 569, 590, 591, 637, 639
- Schrödinger coherent states, 455, 473, 480
- Schrödinger equation, 305, 323, 355, 373, 474, 1399–1400
- Schrödinger picture, 306, *see* Heisenberg picture
- Schrödinger’s cat, 355, 364, 418–422, 514, 527, 640, *see* measurement problem
- Schrödinger’s problem, 641–644, 651, *see* measurement problem
- Schumacher, B. 585
- Schumacher’s quantum source coding theorem (noiseless channel coding theorem), 578, 586, 589
- Schwarz, J., 1298
- selection criterion for representations, 502, 781f.
- self-adjointness, 403, 1400–1402, 1421–1423
- semiclassical analysis, 476
- semiclassical regime, 471
- semidirect product Poisson structure, 461
- Shannon, C., 557, 558
- Shannon entropy, 557, 558, 562, 564, 576, 577
- Shannon’s noisy channel coding theorem, 562, 564
- Shannon’s source coding theorem (noiseless channel coding theorem), 557, 560, 562, 577
- Shor, P., 557, 617
- Shor’s algorithm, 617, 626
- signaling, 373, 391, 392, 582, *see* remote steering
- Simon, D., 617, 623
- Simon’s algorithm, 617, 623
- simultaneity in relativity theory, 237, 248, 256–263
- singlet state, 349, 381, 384
- singularity theorems, 1190, 1207
- size of the visible universe, 1220
- small universe, 1203, 1206, 1230, 1231
- Smolin, J., 575
- Smolin, L., 1293
- Soler’s theorem, 312

- Sorkin, R., 1296
- source evolution, 1221
- space, public and private in relativity theory, 249–252
- space of states, *see* state space
- spacetime
- classical, 263–270
 - full Newtonian, 1387
 - Leibnizian, 1376
 - Maxwellian, 1377
 - neo-Newtonian, 263–270, 1377
 - relativistic, 230–249, *see* Friedmann-Robertson-Walker model *and* causal properties of
 - Gödel spacetime, 1407–1409
 - Malament-Hogarth spacetimes, 1420–1421
 - Minkowski spacetime, 256–263, 1394, 1396–1398
 - Misner spacetime, 1417
 - negative mass, 1422
 - Reissner-Nordström, 1421
 - relativistic, *see* causal properties of
 - Schwarzschild spacetime, 1422
 - Taub-NUT spacetime, 1417
- spatial coordinates contrasted with body coordinates, 54, 76
- spatial homogeneity of universe, 1185, 1209, 1218, 1223–1226
- spatial isotropy of universe, 1185, 1195, 1223
- special initial conditions, *see* fine tuning
- specific heat of solids, 1080
- spectral action, 1296
- spectral family, 278, *see* PVM
- spectral modal interpretations, 375, *see* modal interpretations of quantum theory
- spectral theorem, 278, *see* PVM
- spectral triple, 1296
- spectrum of the area, 1301
- spherically symmetric universe models, 1223
- spin, 291–298, 320, 336, 375, 450, 462
- spin- $\frac{1}{2}$, 290, 402
- spin-network, 1301
- ‘square root of NOT’ gate, 616
- semiclassical analysis, 472
- STC* (symmetric tensor category), 817, 853, 870, 872, 881–888, 896–900, 905–907, 910–913, 919, 920
- STR (special theory of relativity), 154, 164, 188, 256–263, 397, 1347, 1393–1399
- stable causality, 1370, *see* causal properties of relativistic spacetimes
- standard candles, 1200, 1221
- standard model of particle physics, 683–723, 1287
- standard models of cosmology, 1188, 1223
- start of the universe, *see* big bang
- state space in quantum mechanics, 280–282, 284, 289, 300, 308, 313–322, *see* mixed state, KMS states, state space
- Steane, A., 630, 631, 650
- Stein, H., 258, 259
- Stern–Gerlach device, 291, 385
- Stinespring’s theorem, 456
- stochastic dynamics, 1037–1062
- approach to equilibrium and objections, 1046–1056
 - ergodic theorem for Markov processes, 1049
 - interventionism, 1037, 1054–1057
 - Markov processes and irreversibility, 1038–1039, 1057–1062
 - Markov processes defined, 1039–1043
- Stone’s theorem, 301, 304
- Stone-von Neumann theorem, 342, 447, 450, 1423
- stress-energy tensor, 1394, 1410, *see*

- energy conditions
- strong energy condition, 1187, 1211
- strong non-disturbance, 353, 354
- structure constants, 37
- structure formation, 1201, 1202, 1207, 1209, 1210, 1214, 1217, 1221, 1228, 1244, 1245
- structure function, 86
- subadditivity inequality, 578
- subcategory, 864
 - monoidal, 865
- submanifold, 40
- submersion, 40
- substantivalism, 18, 1312, 1411, *see* container view of spacetime
- super
 - Hilbert space, 885
 - fiber functor, 886
 - group, 885
- super-horizon structure of the universe, 1206, *see* observable universe
- supernovae, 1194, 1200, 1224, 1227
- superposition, 280, 281, 317, 318, 356–358, 378
- superselection sector, 501–506, 778f.
- supersymmetry, 724, 1299
- surface brightness, 1197
- symmetric algebra, 911
- symmetric sequence, 494
- symmetry, 252–256, 301, 321, 331, 337–339, 867, *see* gauge symmetry, general covariance, Noether's theorems
 - and group theory, 1337–1342
 - and irrelevance, 1360–1361
 - and Leibniz's Principle of Sufficient Reason (PSR), 1333, 1335
 - and objectivity, 1363
 - arguments, 1332
 - broken, 668–669, 693, 839–840, 1334–1336
 - classification of, 1344–1345
 - continuous, 1345
 - discrete, 1345
 - gauge, 1344, 1345
 - geometric, 1344
 - global, 1360, 1361
 - group, 1342
 - Hamiltonian, 12, 58
 - internal/external, 1344
 - Lagrangian, 161
 - local, 1359, 1361
 - meaning of, 1342–1345
 - of laws, 1332
 - of objects, 1332
 - of solutions, 1332, 1336
 - of states, 1332, 1335, 1336
 - quantum field theory, 764–766
 - principles, 1332
 - transformations, active and passive, 1342
 - variational, 161
- symplectic foliation of Poisson manifolds, 91, 95, *see* symplectic leaf
- symplectic form, 8
- symplectic group, 55
- symplectic leaf, 96, 318, 458, 508
- symplectic manifold, 8, 149–151
- symplectic map, 55
- symplectic vector field, 12
- system of covariance, 329
- system of imprimitivity, 331–334, 410, 447–454, 461–462, *see* Mackey's theorem
 - generalized, 454
 - relation to deformation quantization, 461
- 't Hooft, G., 1292, 1296
- Tannaka, T., 883
- TC* (tensor *-category), 870
- temporal orientability, 1409
- tensor
 - category, 865
 - category, braided, 867
 - category, finitely generated, 900

- category, symmetric, 867
 - functor, 866
 - subcategory, 865
 - thermodynamics, 925, 932–941, 994,
 - see* thermodynamic limit
 - thermodynamic limit, 1017–1028, 1159–1160
 - Thomas, R., 221
 - Thomas–Fermi model, 1085
 - time homogeneity, 304
 - time of decoupling, 1191
 - time-dependent systems, 191–196
 - time-energy uncertainty relation, 344
 - Timpson, C., 632
 - Tomita–Takesaki theory, 735–738, *see* KMS states
 - trace, 871
 - transformation group C^* -algebra, 452
 - transformation theory, 1345–1346
 - transition probability, 318–321, 479, 502
 - classical limit, 473
 - transition probability space, 319, 320, 503
 - transport equation, 484
 - tunnelling, 1156–1157, 1402
 - Turing, A., 614
 - Turing computability, 1418, *see* Turing machine
 - Turing machine, 614
 - nondeterministic (probabilistic), 615
 - universal, 556, 614, 1392
 - twist, 872
 - two-slit experiment, 347, *see* double slit experiment
 - typical sequence, 558, 559
- Umdeutung*
- functorial quantization, 471
 - Heisenberg, 427
 - of classical pure states, 474
 - quantization, 446
 - Rieffel, 459
- Schrödinger, 428
 - von Neumann, 430
 - unbounded energies, 1232
 - uncertainty relation, 341–349, 1289
 - algebraic derivation of, 344
 - optical derivation of, 343
 - wavefunctional derivation of, 343
 - uniform thermal histories, 1225
 - uniqueness of the universe, 1216, 1219
 - uniqueness of cosmology, 1232
 - unitarily equivalent representations, 447, 732, 843, *see* representation
 - unitarity of a Hamiltonian flow, 509
 - unitary representation, 301
 - Unruh effect, 1143, 1295, 1303, 1424, *see* Hawking effect
 - Unruh, W., 1295
- Vaidman, L., 597
 - van Fraassen, B., 647
 - van Hove limit, 1135, 1160–1162
 - Veltman, F., 1292
 - Vidal, G., 631
 - von Neumann, J., 430, 432
 - von Neumann algebra, 443, 498, 506, 522, 730, 731, 742, 744–756, 762–766, 782, 802, 806–808, 812, 839, 1101–1105, 1110–1114, 1118–1121, 1126, 1144, 1146, 1148, 1163, *see* C^* -algebra
 - hyperfinite, 569
 - type classification, 733–738
 - von Neumann chain, 432, 515
 - von Neumann entropy, 575, 577, 579, 1091
 - von Neumann’s imprimitivity theorem, 450
 - von Neumann’s infinite tensor product, 504
 - von Plato, J., 962, 967
- Wallace, D., 646
 - wave packet revival, 429, 519

wave-particle duality, 345,
wavefunction, 281, 343, 373, 379–381
weak non-disturbance, 352, 354, *see*
state in quantum mechanics
Weber, T., 645
Weyl algebra, 758
Weyl quantization, 460, 473
Weyl relations, 333, 340–3411
Weyl, H., 233, 234, 1340, 1363, 1369
Wheeler, J., 650, 1291
Wheeler-DeWitt equation, 1293
Wien’s law for blackbody radiation,
423, 1078
Wiesner, S., 556, 593
Wigner function, 479
classical limit, 480
Wigner’s friend, 377, 1164, 1166
Wigner’s theorem, 300
Wigner, E. P., 1362
and his hierarchy, 1359–1360
geometric symmetries, 1344
Wilson, K., 1300
Wittgenstein, L., 514
WKB approximation, 484
Wootters, W., 580, 590

yes–no experiments, 310

Zeilinger, A., 632
Zurek, W., 580

This page intentionally left blank